The Cognitive Life of Maps

Roberto Casati

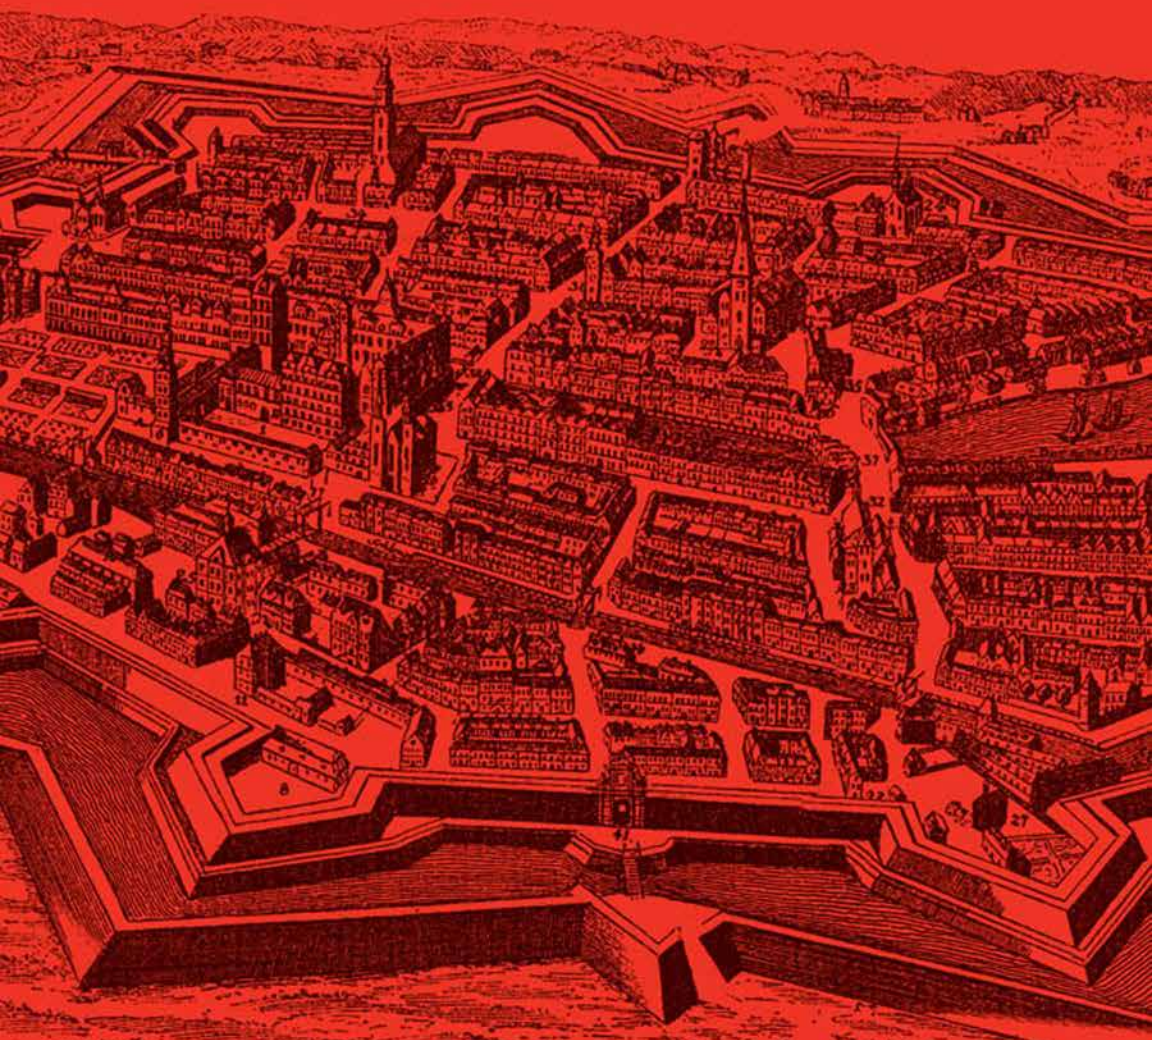

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

The Cognitive Life of Maps

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

The Cognitive Life of Maps

Roberto Casati

The MIT Press
Cambridge, Massachusetts
London, England

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

© 2024 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-ND-NC license.

Subject to such license, all rights are reserved.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Casati, Roberto, 1961– author.
Title: The cognitive life of maps / Roberto Casati.
Description: Cambridge, Massachusetts : The MIT Press, [2024] | Includes bibliographical references and index.
Identifiers: LCCN 2023033657 (print) | LCCN 2023033658 (ebook) | ISBN 9780262547086 (paperback) | ISBN 9780262377232 (epub) | ISBN 9780262377225 (pdf)
Subjects: LCSH: Space and time. | Maps. | Cartography.
Classification: LCC BD621 .C33 2024 (print) | LCC BD621 (ebook) | DDC 114—dc23/eng/20231117
LC record available at https://lccn.loc.gov/2023033657
LC ebook record available at https://lccn.loc.gov/2023033658

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Contents

Acknowledgments ix

An Overview 1
1 A Toolbox and the Mechanics 11
2 The Central Features of Maps: Systematicity, Hypergenerativity 41
3 Pictures and Maps, Pictures as Maps, and Some Other Maps 83
4 Maps, Navigation, and Situation: Red Dots, Green Dots 109
5 Maps and Models 133
6 Clockfaces 151
7 Music Notation as Map 167
8 Organizers as Maps 181
9 Concluding Remarks: The Cognitive Life of Maps 201

Appendix: Maps in the Brain? 213
Notes 217
References 225
Index 239

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

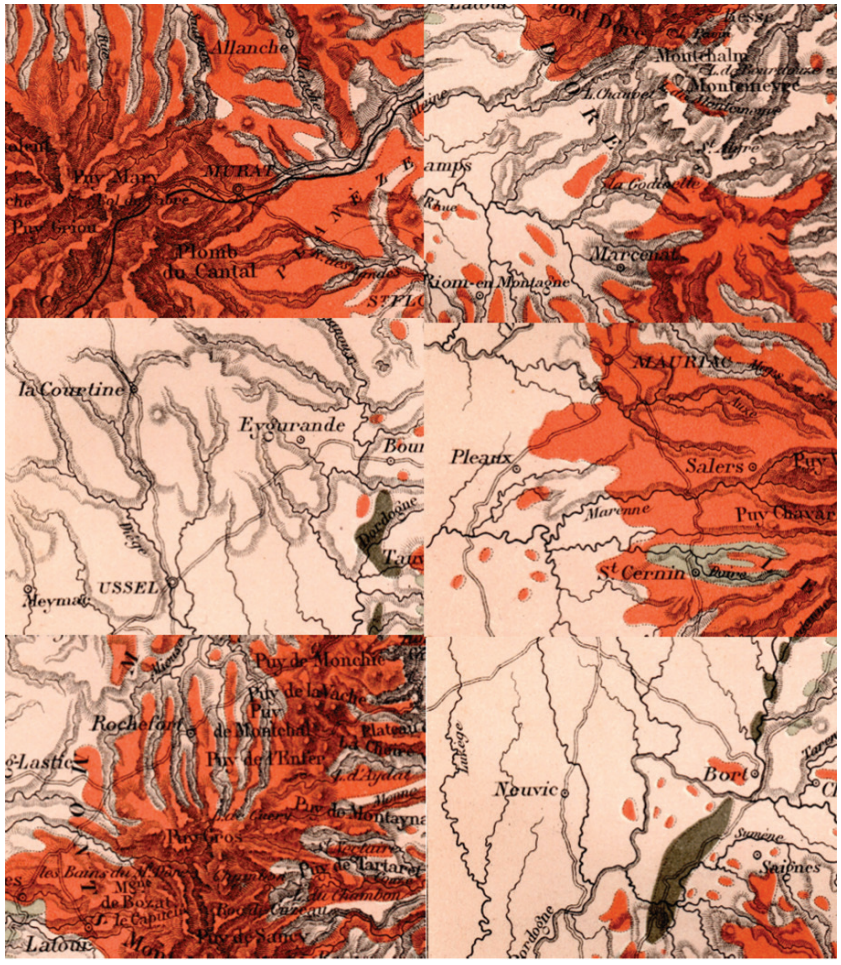

Is this a map? Each of its six pieces are rather uncontroversial maps of six distinct territories. I cut a map into six parts and reshuffled the six parts. Thus, the sum of the six represented territories is a single region. But is the figure above a map of that region? What makes a map a map? Author's collection. Image credit: RC, from *Nouvelle Geographie Universelle*, Paris: Hachette, 1877.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Acknowledgments

Many, many thanks to all the people I exchanged views with, who inspired me in many ways and who allowed me to err in ways they may not endorse. Barbara Tversky is at the top of my list; everyone else will understand. Thanks to Catherine Abell, Mariela Aguilera, Katerina Bantinaki, Axel Barcelo, Massimo Bergamasco, Quentin Blomet, Luca Bonatti, the late Valentino Braitenberg, Denis Bühler, Marco Buiatti, Elizabeth Camp, Barbara Carnevali, the late Richard Carter, Pompeo Casati, Patrick Cavanagh, Robin Champenois, Paul Chu, Michel Denis, the late Rino Desio, Jérôme Dokic, Jacques Droulez, Pablo Fernández-Velasco, Eloï Fiquet, Andrea Formica, Matteo Fossi (for some extreme page turning), the late Christian Freksa, Nivedita Gangopadhyay, Valeria Giardino, Maurizio Giri, the late Vittorio Girotto, Sergio Giudici, Richard Glauser, Valentina Gliozzi, the late José-Luis Guijarro, Ed Hutchins, Jenann Ismael, Pierre Jacob, LaDona Knigge, Robert Koester, Jeremy Kuhn, John Kulvicki, Glen Lomax, Diego Marconi, Nirmalangshu Mukherji, Marie-Vic Ouzuf-Marignier, David Mark, Farid Masrour, Kevin Mulligan, Gloria Origgi, Elena Pasquinelli, Manuela Piazza, Christopher Peacocke, Bastien Perroy, Francesco Pierini, Goffredo Puccetti, Kapil Raj, Jonathan Raper, Anne Reboul, François Recanati, Paulo Santos, Marcella Schmidt di Friedberg, Gianfranco Soldati, Dan Sperber, Stéphanie Roisin, Johannes Schöning, Sevil Şeten, Barry Smith, Magda Stanová, the late John Stewart, Frédérique de Vignemont, Alberto Voltolini, Ken Walden, Leo Zaibert, John Zeimbekis. And, of course, Achille Varzi.

Thanks to my navigation teachers. Thanks to the countless students and attendees of many classes, courses, seminars, and workshops: at SUNY Buffalo where I taught in the framework of an IGERT Grant at NCGIA, at NYUAD, at Columbia, at ENS, Paris, at EHESS, Paris, and in particular the *Maps Semantics* workshops, the *Maps, Diagrams, Sketches* seminar (run with

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Valeria Giardino), and the *Disorientation Seminar* (run with Bastien Perroy and Francesco Pierini).

This work has received direct and indirect support from grants ANR-10-LABX-0087 IEC, ANR-10-IDEX-0001–02 PSL, Disorientation Remediation (PSL-NYU), Map Semantics (PSL-Columbia), Knowledge in the Wind (EHESS).

Earlier versions of the chapters on clocks (Casati 1997), music notation (Casati 2016), Venice maps (Casati et al. 2017), and a larger discussion of numerals (Casati 2014) have been published previously and have undergone significant remodeling in the light of the general view advocated here. I thank the editors of the respective relevant edited volumes, Leo Zaibert and Anne Reboul, for granting permission to reuse some of that material. I also thank John Kulvicki and John Zeimbekis, with whom I coauthored the text on Borgesian maps, and the editors of *Analytic Philosophy*, as well Magda Stanová, with whom I coauthored the text on Venice, with the assistance of Stéphanie Roisin.

The writing got started during a research visit at the Leslie Center for the Humanities at Dartmouth College in 2013. The bulk of the book was written during a research visit at the Italian Academy of Columbia University in 2019. There are no places like these for getting work done in such ideally supportive conditions. I would like to thank the staff of these residencies for their tireless problem-solving, the fellows of the 2013 and 2019 residencies for their critical insights, and then specially Colleen Glenney Boggs at Dartmouth and Abigail Asher, Barbara Faedda, and (last but not least) David Freedberg at Columbia for believing in intellectual free wandering. Many thanks to Frédérique de Vignemont for insisting that I turn the manuscript into a book. Many small events have punctuated the writing, and I am particularly glad to underline the advantages of being sometimes on stage and show how page-turning rescued me from a Catherine Elgin objection about music scores. Reviewers for the MIT Press helped enormously with their detailed comments. Emilie Flamme made my English a bit more idiomatic. A long round of revision was done when secluded during the 2020 pandemic. Taking one's whole family twice to another continent, and then into repeated lockdowns, requires a lot of remapping—not only geographic. Thanks to Nina, Anni, and Lise, thanks to the welcoming school systems in New Hampshire and NYC, thanks to the amazing French school system that kept them intelligently busy during the pandemic, and thanks to Beatrice Biagini for making this book, and much more, just possible.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

An Overview

First of all, the title: the *cognitive* life of maps versus the life of *cognitive maps* versus the cognitive *life* of maps.

The cognitive life of maps is distinct from the life of cognitive maps. In fact, the book (almost) does not deal with cognitive maps at all. And yet, there are two senses of "the cognitive life of maps." Following theories of extended cognition, one may claim that maps have a cognitive life of their own. However, the book follows a different path, the study of cognitive engagement with maps. In a sense, maps are temporarily alive for those who design, draw, and use them. How can they be alive? What kind of life do they have?

The book is organized in the following way. Chapter 1 presents what I consider to be the essential toolbox for studying cognitive artifacts and their mechanics. Here, I briefly rehearse topics from cognitive science and the philosophy of mind. Chapter 2 presents the main claims about what maps are and how they work—their specific syntax, their peculiar semantics, and their pragmatics. The subsequent chapters (3, 4, and 5) delve into the mechanics of maps as they are used for navigation, the differences and similarities between maps and pictures, and the differences and similarities between maps and models. Then, the final chapters (6, 7, and 8) test the skeletal theory of maps on an enlargement of scope: Can we find map-like structures in other cognitive artifacts, and how are these structures specifically map-like? These chapters make the case that clockfaces, music notation, writing, schedule organizers (such as calendars), and numeral series are or contain map structures. This enlargement strategy makes a case for the centrality of maps. Missing an understanding of maps, we cannot understand how a number of other cognitive artifacts work. The

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

book closes with a brief discussion of territorialization and of mental and brain maps.

The book starts (chapter 1) with an unsystematic inventory of some relevant tools from philosophy and cognitive science. It does not purport to cover all the debates on maps or cognitive artifacts developed within these and neighboring disciplines: rather opportunistically, it picks some ideas and methods that can help probe the depth of the cognitive foundations of cartography. I find myself relatively comfortable in a relatively conservative computational-representational framework. I am relatively less comfortable with some newer trends, variously labeled as "extended" mind/cognition, which, however, have the interest, for the present project, of discussing the use of cognitive artifacts (including maps) in (some) detail. I'll point right away to an important non sequitur: from the fact that we can and do "offload" some of our internal cognitive tasks to maps (and other cognitive artifacts), it does not follow that maps (and other artifacts) perform any cognitive operation in our stead. But other interesting things follow that define the mechanics of a post-offloading theory of cognitive artifacts. Based on insight from "dual" or "two-system" theories of cognition, and extending the metaphors they incubate to define not two but four modes (M) of functioning of the mind, we can single out map use from various ways of finding our way in the environment: how you go back to the bodega next door (Mode 1); how you reason out loud and slowly (Mode 2) to instruct someone to reach the intersection of 5th Ave and 107th Street; how you locate and go to Hoboken, New Jersey, using a map (Mode 3); how the GPS takes you everywhere without you thinking of your path (Mode 4). We will distinguish core computations and auxiliary computations relative to a certain cognitive task. This will then allow us to reframe the dialectical bite of "extended" mind theories of cognitive artifacts and will force us back to the drawing board to sketch a computational account of map use. Some other notions, such as that of a representational epiphany, will be introduced.

Readers interested in the account of maps can jump directly to chapter 2. The key constraint of this account is that of map use: What is a map that a mind like ours can use to grasp a territory and navigate it? We first distinguish between types of inferences we can draw from maps: nonsituated inferences ("Since I can see that Sardinia is smaller than France and that France is a subpart of Europe, I conclude that Sardinia is smaller than Europe") and situated inferences ("If *I* want to reach Antananarivo from

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

here, I have to cross *this* part of the sea" or "Chances of rain in the area *I am in* will be 70 percent one hour from now").

We will need some characterizations of maps. These characterizations are provided by a discussion of *contrast classes* that include diagrams, written language, and pictures. We'll also need to distinguish maps from (mathematical) *mappings* and different senses in which maps are *projections*. But contrast classes, as we shall see, are not exactly sharp; they only highlight tendencies and dimensions in a rich conceptual field.

It is important not to get intimidated by differences between representational genera while at the same time paying attention to them. Indeed, along the way of any discussion of cartographic reason, we encounter the bewildering variety of methods used for representing territories by using two-dimensional sheets (of paper or on screens). We should pay respect to this variety without, at the same time, letting it discourage us and push us toward an *atheoretical* stance. Yes, we know there are different types of cartographic projections. Maps can lie and distort. We have contrasting intuitions about east-west or north-south reversed maps, and so on. I propose to distinguish here, and elsewhere in this book, between (substantive) *objections* and (innocent) *Distractions*. To do so, I prompt where a line must be drawn by signaling what I take to be a Distraction in a sans-serif font.

This methodology pays off as I start making some of the central claims. First, I assume a difference between and mutual independence of map and mapped *territory*. Then, I endorse the principle of same-dimensionality of map and territory. Many Distractions will pop up in no time such as "There would be no Italy had there not been a map drawing the boundaries of Italy" and "The map is flat, the mountain is not," and I propose, methodologically, to focus on two simple cases, the 2D map of a 2D surface and the 3D cake (cuttable) model of 3D geological strata, then build from these cases and enrich the account.

Given the dimensionality constraint, a territory must be a 2D surface if it is to be mapped onto a 2D map (not necessarily a *flat* surface and not necessarily a horizontal surface: there are vertical maps of vertical territories, such as geological cross-sections, a vertical cut of an airplane indicating emergency exits, or an elevation map of the Tour de France). This means that territories are, in a harmless sense, abstract entities in the eyes of mapmakers. They are what can be mapped of a certain world region. Maps see territories as layers of abstractions. The physicality of territories is also a

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

methodological constraint. Territories *both* are physical *and* can be abstract. The two features are not incompatible. A donut's center of mass is a respectable physical quantity, and yet it is immaterial because it resides outside the physical body of the donut. Some territories are more easily mappable than others. A minimalist approach to mappability is that the map does not create its territory. Rather, a map draws attention to some aspects of a territory's two-dimensional structures. Maps only see the 2D aspects of a territory. (We will discuss the territorialization of maps in chapter 9.)

Another key claim, following Casati and Varzi (1999), is that it is possible to provide a working, if skeletal, map semantics by translating maps into their formal equivalent—*formal maps*—to elicit and account for *some* map-based inferences (in particular, both situated/navigational and non-situated). Formal maps are the normal form of maps: they decompose into atomic map stages or layers, coloring is akin to predication, regions are akin to singular terms, and evaluation is done by evaluation of atomic layers and superposition of layers. The translation of ordinary, everyday maps into their *formal* counterparts requires some ingenuity but is worth the effort as it allows focusing on some key syntactic and semantic features of maps in view of their use.

Here, of course, problems and Distractions abound: Will there be a sufficient color supply? Can the syntax of formal maps really support superposition of regions? Do formal maps deliver visual clogging? What about the four-color theorem? What properties can colors express? How do we cope with the oddity of map regions construed as structured singular terms? At a minimum, they all point to the limits of the expressive power of formal maps—limits that in some sense are to be expected (as they are expected with other formalisms). More serious problems go under the title of the absence intuition, the expressibility of negation, and the nature of predication.

Using formal maps as the basis of the discussion, further claims concern:

- the holistic automaticity of assignment of parts of the territory to parts of the map. This means that all maps, including empty maps, have semantically active parts. (Here is a tiny outline map of the room I'm in right now: □. Its parts refer to precise parts of the room, although it may be hard to tell precisely which part refers to which.)
- the hypergenerativity of maps, building on the "natural generativity" of pictures (Schier 1986). Consider: Knowledge of the meaning of "My cat attacked your dog" and perceptual acquaintance with dolls, smiles, and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

cars does *not* generate knowledge of the meaning of "My doll smiled at your car" (particularly if you don't know these English words). But knowledge of the meaning of a picture representing a cat attacking your dog and perceptual acquaintance with dolls, smiles, and cars does generate (some) knowledge of the meaning of a picture of my doll smiling at your car. The mechanisms of the natural generativity of pictures are open to empirical investigation, although we appear to understand their general architecture. In the case of maps, it turns out that maps *hypergenerate*. As soon as physical tokens are placed in space and assigned some spatial reference, they can be taken as representations of spatial relations between their references. Perceptual acquaintance with the territory is not needed; spatial competence alone suffices. Map generativity is thus specific. Which is the underlying mechanism? We'll start from a nexus. First, maps are highly abstractable to the point of easily reducing to formal maps. Second, their hypergenerativity flows from knowledge of the automaticity of the structured assignment of territory regions to map regions.

From the same-dimensionality of map and territory, another important nexus emerges and takes us to a discussion of the relationship between maps and pictures, the visuality of maps, and the perspectivality of pictures, to consider some contrast facts (e.g., maps are poor at representing occlusions, contrary to pictures). This is the subject matter of chapter 3. There is an all-too-strong pull toward considering maps as a kind of picture or vice versa. The interchangeability of *map view* and *aerial/satellite view* in geographic software counts perhaps as a motivation for the pull. Another motivation is the *picturification* of maps to make them more easily interpretable visually. In a sense, maps and pictures are cousins.

I propose a simple theory. Pictures *are* indeed maps but maps of a very special kind. They are maps of the Albertian Window used to visually frame the picture. Sure enough, chapter 3 goes into some depth to consider a pile of Distractions: the difference between flat and two-dimensional realistic pictures versus abstract pictures, and photorealistic pictures of unicorns. Above all, we should mind a non sequitur: from the fact that pictures of an X are maps, it does not follow that they are maps of X. We should also mind the converse non sequitur: from the fact that pictures map the Albertian Window, it does not follow that they *depict* it. Here we find the possible resolution for how features of "sense data," "views," "aspects," "looks,"

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

"outline shape," and "occlusion shape" account for pictures. The table is round (measurably so). But on the canvas, to represent the table, an oval is found (measurably so). But pictures do not *depict* visual angles (aspects, looks, occlusions). Rather, they *map* those items. Both pictures and maps (and some diagrams too) are examples of analog representations. We shall point out some distinctions, in particular to account for the fact that an analog representation of a spatial territory, using space as a representational vehicle, and from which spatial inferences about that territory can be drawn, may not be a map.

How can maps be used? GPS-based navigation devices have familiarized us with dots moving on maps, representing our position. Chapter 4 discusses the difference between tracking and situated tracking, between what I shall call Green Dots and Red Dots. A webcam on my handheld GPS-based device generates a Green Dot image that is broadcast onto all sorts of different places. The Green Dot tracks me but can only be used for navigation when it is situated where I am, thereby becoming a Red Dot. The discussion of Green Dots and Red Dots will take us to some of the intrinsic limits of situated cognition accounts. Situation is a sufficient but not necessary condition for Red Dotting; epistemic *segregation* is both a necessary and sufficient condition. Red Dots are then Green Dots with added epistemic segregation.

We will see that maps are *conceptually* secondary relative to situated GPS-based navigation devices—they are abstractions out of GPS-based navigation devices. The cognitive advantage of maps is their *a-perspectivality*, their *orientability*, as opposed to pictures of a landscape, which are oriented and define a canonical viewpoint. We shall also see that all maps are conceptually (not historically) tributary of a very peculiar type of maps, Borgesian maps (that is, 1:1 superposed mappy representations of a territory). Borgesian maps have oftentimes been declared absurd or irrelevant, but their existence and role will be vindicated. Not only are there plenty of them but they are crucial in understanding abstraction and territorialization.

The rest of the book is an exploration of map-like structures. Having worked with a sufficiently stable theoretical notion of map, based on the same methodological restrictions, will pay off here as we can enlarge the field of phenomena under study. Some structures are not generally labeled as maps: radar "images," fMRI "images," and nano "images." Some, such as nano-maps, are picturified unwillingly. Others, such as fMRI, are explicitly picturified. Signs of *picturification* include added shadows and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

occlusions—indeed, as shadows are traces of occlusions. You may find yourself seeing maps in many cognitive structures: contiguous blocks of numerals, "organizer" sequences, calendars, music scores, clockfaces, tables of contents, written language, people in a line, star patterns. Some require particular care. Within each structure, we will focus on the exact delimitation of their respective territories. For instance, a clock is a device for Red Dotting a calendar, which itself is a map in Mode 4. In the past, we did a Red Dot on vintage tear-off calendars in Mode 3.

There will be examples of unusual cartographies. I am not a cartographer, but my old, predigital minor in visual design and decades of involvement in cognitive approaches to education give me motivation for a hands-on approach to maps. How can we enhance their contribution to understanding certain territories? So, you will find a new cartography of Venice, a redesign of instructional globes, and an app that shows that clockfaces hide maps of the world.

This book is the result of a long series of investigations, attempts, and explorations dating back to the 1990s. Epistemic artifacts have been an interest of mine, and even an obsession, for quite a long time now. How is it that we empower our perceptual and cognitive faculties through interactions with tools of so many different formats and ways of functioning such as paper maps, waiting lines, drawings, and mechanical clocks? What exactly is this empowerment, and how is it brought out as a result of cognitive engagement? What are the mechanics of this cognitive engagement, and do these mechanics have general elements that apply to the whole range of artifacts? What happens to a mind that is equipped with epistemic artifacts? What does it mean to claim that the mind gets "augmented" or "extended" through the use of artifacts, in particular as some artifacts mechanics are meant to *limit* mental operations in order to achieve certain pragmatic goals? Collectively, these inquiries are salient when one starts to discuss maps.

Maps are generated and used billions of times a day. Despite this fact, one hears claims that the map era will soon be over. Maps, as visual cognitive tools, have flourished in the last centuries (the Cartographic Age), but we are told that they will likely be obsolete in the near future. There will still be maps or, better, data structures that can occasionally generate visualizations resembling what we today call a map. However, the very idea of a graphic display used for spatial orientation and specifically a *support* for wayfinding will be replaced by proxies of a different nature: *deferential wayfinders*.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

It is already the case that the most used instrument for wayfinding is a GPS-based navigation system with a voice interface that tells millions of drivers every day what action they are expected to perform next to get to their planned destination (Bray 2014). The system does not actually need to display any map (and safety concerns about distracted driving may even suggest that the navigation system should be prevented from displaying anything visual). Displays already shifted from "north-up" to "ahead-up" to ease browsing. It should be noted that I am not in the business to assess the debate about the cognitive loss related to intensive use of cognitive proxies that some lament (Carr 2014; Lin et al. 2017). Possibly this loss is completely offset by gains that, again and again, automation makes possible. The Costa Concordia accident of 2012 (Marine Casualties Investigative Body 2013) and the 2014 Team Vestas accident during the Volvo Ocean Race (Oxenbould et al. 2015), among many others, warn us about the limitations of the human ability to govern large and fast vessels unassisted or inadequately assisted by navigation technologies.

So why study maps? And what can we learn from their study? One reason could be just documentary. Before map use becomes obsolete, we may want to be able to document what has been for centuries our interaction with a crucial cognitive artifact. Some map users have not yet initiated, let alone completed, the transition to proxy navigational instruments; these individuals would be the privileged target of a study in, say, cognitive anthropology. Ed Hutchins did invaluable documentary work when he investigated complex navigational techniques on military vessels that have since been superseded by more technological, GPS-based options. In this work, Hutchins (1995) unearthed the phenomenon of distributed cognition. Generally, we should consider the study of maps as a subdivision of the study of cognitive artifacts. Maps are particularly interesting in this respect because their relative simplicity hides a complex set of interactions that involves various branches of the cognitive architecture deployed when we use them. And then, perhaps, we should not be so pessimistic about the future of maps. Maybe the *navigational* era of map is over, but other *non-navigational* uses will be with us for a long time. Representing space *by space* is an important, cognitively convenient, and, in the end, irreplaceable way of representing space. In so many circumstances of our educational life, our business, and our political decision-making, the conversation starts with a map, relies on maps, and ends with a map.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

But there is more. Many reviewers of the manuscript of this book have suggested that I should focus more on the political aspects of mapmaking and use. Now, as a person involved in issues of ocean conservation, I am very much aware of the enormous power of maps in framing political conversations. The definition of Exclusive Economic Zones and of Marine Protected Areas could be seen as immensely problematic case studies for the cognitive tendency to project, through mapping strategies, land-based categories onto the fluid ontologies of the ocean, whose biology, chemistry, physics, and human uses are inherently recalcitrant to the drawing of boundaries and the definition of "territories." Sure enough, the main thrust of the theoretically minded literature in geography in the last half century has concerned the issue of the conceptual ties between maps and territories, justly warning about the role of maps in defining territories and in projecting power patterns over territories (Casey 2002; Fainstein and Campbell 2016; Robinson and Petchenik 1976; Harley and Andrews 2002; Wood and Fels 1992; W. Rankin 2018; Varzi 2001; Winther 2020). I will say something about these in chapter 9. My timid concern with most of this literature is that it is—surprisingly—not particularly sensitive to maps' *specificity*. For instance, the definiens "a graphic representation of the milieu" (Robinson and Petchenik 1976) applies to maps but also to diagrams, photographs, and even written texts. How do we distinguish between them? What difference do maps have that make them so powerful, as opposed to texts and diagrams, themselves instruments of power? What is the added value of maps, their spatial rhetoric? The present book can be seen as an attempt at answering these questions. There is an intriguing passage in a novel by Jim Harrison (1994) on the irrelevance of the ocean for mountain communities when they learn that they live at some altitude above sea level. "One Stab disliked the altitude numbers attached to mountains. How high above which of the seven seas Tristan had told him about? What did the numbers mean if there were no sea near them?" If the enormous, uncontrolled power of maps is in their making distant places relevant to local realities, and in the end, in making each place relevant to all other places, we need to explore the mechanics of the semantic cartographic engines capable to make us think obsessively in terms of connections between places. And if a map-free society is today a chimera, we can at least try to prevent a map-ruled society from being a nightmare.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

1 A Toolbox and the Mechanics

This chapter introduces a number of theoretical tools that will guide our exploration of maps. It is neither a full-blown theoretical introduction nor a state-of-the-art review or methodological manifesto. Rather, this chapter is a toolbox of ideas, whose list can let the reader get a sense of the type of investigation undertaken here. Readers may skip this chapter and get back to it when they feel they need to double-check some of my theoretical prejudices. However, the interest of these tools is precisely that they allow us to see commonalities and differences in the vast set of graphic productions, of which maps are a subset.

Computation and Representation

> Navigation is computation.
> —Edwin Hutchins, *Cognition in the Wild*

First of all, this book works within a metatheory, a framework derived from authors such as Fodor (1983) and Marr (1982), among others. The maps discussed here are produced and used by people with a mind-brain (other options are possible: maps produced and/or used by artificial systems). The metatheory considers the mind-brain as an information processing engine that manipulates information gathered through sensory organs to computationally build internal representations of the environment to further guide action and make thought possible. The representational-computational account is implicitly taken for granted by current psychological research, as the account makes it possible to perform a simple major intellectual step, distinguishing the contributions of the environment from those of the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

computing machinery. Every time psychologists deal with illusions (such as the Müller-Lyer illusion) or biases (such as the confirmation bias), they rely on some normative elements that are packaged with the notion of representation: a representation can be true or false, adequate or inadequate. Competing metatheories, such as those that rely on less normative notions like mirroring, handle illusions and biases poorly. Illusions and biases, in turn, fuel psychological research. Each time an illusion is discovered, it is set to be explained in terms of the contribution of the computing system to the final representation. Working with representations and computations allows further decomposition of the process into theoretically significant steps, expressed in terms of rules. This in turn can eventually lead to the formulation of some assessable predictions. This process uncovers a cognitive architecture, a way in which components of the system are functionally related to other components. Finally, the paradigm favors the integration of cognitive hypotheses with studies of the underlying biology, of the ecology of relevant environments, of individual development, of impairments, and of evolutionary constraints. Importantly, the internal representations are theoretical posits whose properties are the target of empirical research.[1]

Ways of Functioning of the Mind/Brain

The psychological literature endorses (although nonconsensually) a distinction between two "systems," or two modes of operation of the brain in certain tasks (Kahneman 2013). These are mainly reasoning and decision-making tasks. Mode 1 (M1) is described as an automatic, fast, and fluent operating mode that delivers rough but locally acceptable results. Mode 2 is described as a mode regulated by will and attention that operates slowly and stepwise, is taxing to memory, and delivers comparatively more accurate results than M1 in many cases. An example from spatial navigation is: you can find the exit of the library that you entered for the first time ten minutes ago without having to consult a GPS-powered device, look at a map, or ask for directions. In this case, your spatial brain does all the work for you, fluently, in Mode 1. If, on the other hand, you want to go from Washington Square to Columbus Circle in New York City, you may have to think step by step, to plan, to keep waypoints in memory, to deploy attention. This is a Mode 2 way of operating. There are nowadays a number of ways to present the distinction between M1 and M2 and to relate it

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

to other older distinctions—intuition versus reason, heuristics and biases versus logic, encapsulated modules versus central processing (Evans 2003; 2012). The simplest, least controversial characterization of M2 is that in M2 there is a high demand on attention and working memory.

Talk of modes makes a first and important generalization possible. There exists a large variety of mental engagements beyond verbal problem-solving and decision-making. As we have just seen, spatial navigation is hospitable to the two-modes (M1 and M2) distinction. Visual analysis is another example: you can just let vision process information automatically or you may want to intervene in it by directing your overt attention to specific details (figure 1.1). Vision is a bag of routines that operate in M1 mode from color constancy to depth via binocular disparity to pop-up mechanisms of visual attention.

Vision, though, can be modulated by attention. You could focus your attention, overtly or covertly, toward details in the scene and remark items that would otherwise escape at first glance. Careful *observation* is to careless vision what careful thinking is to biased, intuitive thinking. With its demands on attention and working memory, careful observation is the M2 of vision. At the same time, observation is subject to the will and yet relies on M1 visual routines.

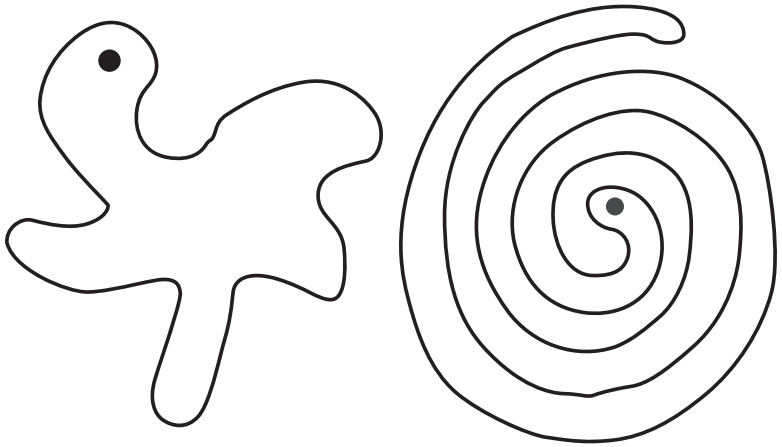

Figure 1.1
Is the dot inside or outside the blob? (Quick answer, visual routine, M1.) Compare: Is the dot inside or outside the closed spiral? (Slow answer, no visual routine, M2.) Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Although careful observation and careless vision are likely to be grounded in different architectures (see Cavanagh [2011] for a description of the attentional mechanics), the main point of the two-modes distinction is that the differences observed are internal to the visual *system*. Therefore, we cannot generalize talk of two *systems* because the relevant commonalities are not there between vision and navigation. But we can generalize talk of two modes.

Cognitive Advantages

We would not use cognitive artifacts such as maps if they did not provide *cognitive advantages*: less burden on memory, enhanced decision-making, eased visual inspection, and faster processing. Cognitive advantages are related to processing costs. These costs can be measured in various ways (e.g., time to accomplish a certain task, effort, fluency, and retrieval).

A simple illustration of trade-offs in cognitive advantages is provided by a comparison of maps and photographs when used for navigation (see chapter 4). *Portability* is both the cognitive advantage and signature of maps over other environmental representations such as photographs. Photographs are, by etiology and design, *viewpoint bound*. One may use a photograph as a guide to explore a certain environment, but the limits of this use become quickly evident precisely because photographs are viewpoint-bound. If you move around, you will find after a while that some important landmarks the photograph represents are out of view and other landmarks will become salient that are not represented in the photograph. Parallax, hence perspective, changes dramatically as you walk in the environment. This means that when using a photograph as an aid to navigation, the inferential load grows exponentially as one moves away from the photograph's viewpoint and becomes more and more framed in a hypothetical format ("If this is Merril's Lodge, *then* that must be the Blue Fountain"). On the other hand, the inferential load required for using a map is constant over the whole territory mapped.[2] The M2 engagement with maps fades into a number of automatic M1 routines as you start navigating, whereas the M1 immediate recognition of a landscape in a photograph turns quickly into M2 inferential reconstructions of the ground as you move away from the photograph's viewpoint. In the design and adoption of cognitive artifacts, there is a constant trade-off of cognitive advantages and disadvantages.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

A representational advantage (Casati 2003b) is a type of cognitive advantage. A representational advantage occurs when, for instance, certain parts of a scene are computed as objects, as opposed to those that are not computed or computed as non-objects. A piece of wood has an advantage over its shadow: we tend to neglect the shadow as we do not see it as an object. The piece of wood, as a figure, has various advantages over its background: it "owns" the visual boundary that separates it from the background.

Mode 4 (Mode 3 Will Be Presented Later)

My proposal is to extend the M1/M2 metaphor/tool to include another mode, M3. Actually, I propose to introduce *two* other operation modes: Mode 3 (M3) and Mode 4 (M4). To understand the importance of the introduction of M3, I find it convenient to start from M4. M4 occupies an extreme end of the spectrum used to represent ways of operating under examination here. Working in M4 means *delegating* (almost) the *whole* relevant mental process to some external *device*, and only caring about the *solution* of the delegated process. M4 really is a proxy mode. For instance, you launch a navigation query ("destination: Eiffel Tower") and take directions from a GPS-based navigation device by listening to the instructions dictated to you in a headset. You hear, "Proceed to the intersection. At the intersection, make a left turn." You do not even look at the map on the display. This is a paradigmatic case of M4 mode of operation.

It is clear that this extreme type of human–machine interaction is only possible in heavily constrained environments, such as the road system, where decisions tend to be discrete as the options are simplified: you rely on an ontology of networked one-dimensional objects (routes), left turns, right turns, and little else, where the only "legal" moves are going ahead, making turns, or turning back. You do not, let's say, head for a 37° bearing.

It is also clear that in most cases, and in this one in particular, you do not give up all of your mental life. It is not as if you asked a cab driver or an autonomous vehicle to take you to the Eiffel Tower and then went to sleep. You are still making decisions: you launch the navigation query, you listen to the verbal instructions and interpret them, you check them against your perception of the environment where conflicts may arise ("Was it the first or the second exit at the intersection?"), and you should still be able to solve them. Yet, you are giving up most of the *relevant* mental life, if compared to

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

a more engaged relationship with, say, a paper map. In particular, you have given up *navigational* computations and outsourced them completely, just as you would outsource them to a taxi driver in an unknown city (Madl et al. 2015). The key aspect of M4 that is important for our discussion is thus the characterization of the relevant mental life. In M4, you only access/perform at the input and/or the output of the computation. You do not run any of the relevant ("core") *navigational* computations in your brain.

Core Computational Tasks versus Auxiliary Computational Tasks

A *core* computational task, relative to a given performance, is a task that is essential to obtain the desired result. In navigation, core computational tasks include dead reckoning and determining a direction. You must be able to compute your local position (know where you are), and (possibly based on dead reckoning) you must be able to compute your direction to a target. This can be achieved either by finding a bearing directly to target or by laying down a sequence of waypoints and then finding a bearing to each of them in sequence. You must be able to compute shortcuts and detours (e.g., you see that the shortest path goes through the boulder, but you know humans cannot go through boulders, so you have to calculate a detour). In M4, the core computational task of navigation is both offloaded to and taken up by the device. You do not need to know where you are, and you do not need to calculate any of the next steps. Still, you must compute something—that is, know how to launch a query, how to interact with the device, and how to interpret the result provided by the device. Relative to navigation, these are *auxiliary* computational tasks. They are not navigational computations. Besides, they are not specific to navigation and are in fact shared by many other tasks (e.g., looking up an encyclopedia entry on your PC).

Therefore, a signature feature of M4 is that you do not need to access any steps of the core computational task that characterizes the performance. The nature of the performance changes: it is no longer "navigation" but "competent use of a navigational instrument."

Mode 3: Assisted, Not Proxied

Mode 3 (M3) is distinct from M1 and M2 in that it *essentially* involves the use of external artifacts and is distinct from M4 insofar as it relates to a

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

particular task where delegating to external artifacts is not unconditional—and definitely does not outsource some *core* computational tasks. As an instance, consider navigating an unknown environment by using an old-fashioned paper map. First, this activity is clearly different from navigating (M1-*like*) in a known environment, for which your brain has constructed—you can't help it—an effective internal representation, a "mental map." The notion of a mental map may itself be debated (see Madl et al. 2015 for a review), but I assume the key idea of this debate is that spatial information about a territory is stored in the brain in a map-like, spatially organized format. The M1-*like* features of mental-map mapping and their use are not controversial. You do not need to consult an external map to find your way back to the elevators in this building: fluency and automaticity are granted here. Second, navigating with a paper map in an unknown environment is visibly different from receiving GPS directions or even setting a vehicle on autopilot, which are prime examples of M4 mode of operation. Assisted by a map in an unknown environment, you are carefully trying with each step to identify landmarks, interpret and assess affordances, and keep an explicit mental record of visited places; you are also inhibiting (Stanovich 2009) some of your powerful M1 routines. For instance, your perception-based intuitions suggest proceeding in a straight line, but the map signals an obstacle that is invisible from your current location and you take a detour. To take another example, your mental mapping system hooks onto whatever visual features are available in the landscape, including clouds and waves, which make for very ephemeral landmarks, and you have to inhibit or discard this hooking (Casati 2019a). You interpret the map propositionally, not just as a depiction of the terrain but as generating a set of verbal instructions or a set of inferential cues that are subject to various types of distortion (Tversky 1981; Tversky and Lee 1998; Denis 2018). All this puts you in an M2-*like* mode of operation. But the physical map remains essential here: interaction with the map modulates all the other activities you engage in. Map-assisted navigation is deliberately thinking of and making a connection between you, the map, and your environment. This connection requires interpretations of both the map (to make quick inferences about what to expect around you) and the environment (to assess the presence of relevant landmarks), as you try to relate one to the other.

As we shall see in the next chapters, a map stores information about a territory in a specific spatially structured way. For our present purposes, it

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

is important to note that a map records information about places that *you may or may not have visited* and in particular about places that *you may not be able to perceptually access in the present moment*. A map used for navigation should be capable of supporting certain specific *queries*. For instance, by asking appropriate questions to the map, you can determine that you are in a certain location and that you should move in a certain direction if you want to reach your final destination. To perform these functions in navigation, a map needs *orienting*. Orienting can be explicit (e.g., when you are able to have a meridian's representation on the map coincide with the represented meridian) or implicit (e.g., you master a set of rules for transforming each direction you establish on the map on a direction in the represented territory and you do so no matter the actual position of the map relative to you and your environment). In both cases, you have established *physical relationships* between the map and the environment where the map is situated. Moreover, this requires a certain set of skills on your part that you must learn or develop. Navigation is computation. You should be able to compute—on an oriented map—your current location by finding the appropriate correspondences between represented landmarks and their representations (e.g., the bearings of two landmarks intersect on the map in a point that approximates your current location).

Are you "delegating" cognition to a map in the same way that you were delegating cognition to the talking GPS navigator? Certainly not. When you interact with a map, you engage in a specific mode of operation; *you* are performing the core navigational computations. To recap, you have to be able to ask questions that the map can provide answers to and be able to find the answers on the map in order to make decisions; you actively engage with the map, and this engagement lets you generate the mental representations that you need to make your decisions and start your motor routines. You are core-computing the position and route yourself. This interaction with maps I take as constitutive of M3.

The alternate view to M3 is that the specificity of the use of cognitive artifacts in cognition is orthogonal to the M1/M2 distinction: one may use a cognitive artifact in Mode 1 or in Mode 2 with no need of a specific "artifact engagement mode." Most cognitive artifacts are likely approached in M2, cautiously, with the deployment of attention and the exercise of working memory as we learn to use them. Then, training with them can make operation with them smooth and bring interactions with them into an M1

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

style. As a general example, I may have had initial hesitations with the use of inverted goggles, periscopes, or the rearview mirrors that cause me to check at different moments what I am precisely doing, but, after some practice, I became fluent (Pacherie 1995). For instance, drivers of right-driving countries experience difficulties with the rearview mirror of a left-steering car (Hecht et al. 2005). This difficulty likely depends on the fact that learning is necessary to use rearview mirrors and that learning is at first performed in M2.

In a deflationary account, one deploys different computations rather than being in different modes. One can learn to use an artifact in M1, or one can keep painfully using an artifact in M2. I take no stance here about the deflationary position, but I find it convenient to single out a class of cognitive performances that essentially involve artifacts.

The Extended Mind versus the Opportunistic Mind versus the Equipped Mind

The previous section is meant to provide a lens on the debate over the metaphor of the "extended" mind that has consumed philosophers and (some) researchers in the cognitive science community over the last twenty years or so (Sprevak 2019). Actually, I would like to be helpful here and offer a way to escape the thorny aspects of this debate altogether. The debate's main tenet is the idea that cognitive artifacts "extend" the mind, thus that the mind is "larger" than the skull and spreads over onto "external" items. This narrative is fascinating because it is somewhat counterintuitive, and thus interesting and rewarding, and has stirred debate because it opposes traditional, more austere views of the mind as located in the body and, in particular, the brain. The latter views are typically charged with accusations of neglecting the extracranial world: the body, the physical environment, and the social environment. A second related notion is that humans "offload" to cognitive artifacts some of the operations that they (their brains) normally perform without helping themselves to artifacts. Put the two claims together and you can definitely derive some excitement from thinking that you have some actual pieces of *your* mind working out there for you in archives, books, pocket calculators, inscriptions, other people's brains, the internet. You will get an extra thrill by attributing a cognitive life to some of these artifacts: the route to San José "knows" where

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

it is going, and if you follow it, you'll see that it does because you will end up in San José (Haugeland 1998).

The extended mind idea is oftentimes paired with other claims. These include that the mind is embodied, that it is embedded, and that it is enactive. There even is an ongoing institutional trend to package all this as the "Four-E" account of the mind. I think that something simple can be said here. In particular, it can be claimed that the mind *is* embodied, or *is* embedded, or *is* enactive, or maybe any combination of these, but it is *not* extended.

There are many mutually independent reasons to work with notions of an embodied, an embedded, and an enactive mind. We shall indeed spend some time discussing a workable notion of embedding for map usage, whereby maps can only be used for navigation if they not only represent the places we want to navigate but also bear nonrepresentational relations to those very places. Thus, the embedding of maps in a place is a key element of navigation. (In chapter 4, however, I will propose a deflationary account of the phenomenon that embedding addresses.) A workable notion of embodiment allows us to make sense of some features of spatial representations—for instance, orientation of the visual scene relative to our particular connection to gravity or the felt difference between left and right. And there is some good in the idea that at least some perception is perception for action, and thus that perceptual representations encode action-related features of the environment such as affordances (a handle that is the size of my hand affords differently than the toy handle on a toy door). Of course, this is not to say that all versions of embodiment, embedding, and enaction are useful or viable. And, it is much less to say that we should accept all of them at once. There are those who find it interesting to see embodiment spread over and explain language and thought, including researchers who claim that perception *is* just simulated action. We could spend a considerable amount of time trying to assess or just make useful sense of these claims.

On the other hand, there are many reasons *not* to work with the notion of an *extended* mind. The most straightforward reason is that there are no ways to negotiate the contours of the putative extended mind in a scientifically useful way.

As a matter of fact, objections to the extended mind narrative have been addressed from many quarters (see, among others, Marconi [2005]; Sprevak [2019]; Di Francesco [2007]; Fodor [2009]). Paper and pencils, notebooks,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

lists, and pocket calculators make for interesting but harmless candidates for the "extension," and some of the powerful rhetoric of the narrative flows from similar examples of cognitive artifacts. But other candidates are uncooperative. Does my mind include that of my sister when she helps me with my homework? Am I part of her mind when I help her with hers? What about us helping each other? What does it *feel* to be part of someone else's mind or to have someone else's mind as part of one's own, or both? When I use the Pole Star for finding my way on a clear night, is *it* a piece of my mind? (Notably, given light speed, the star as I see it may have already ceased to exist.) If I am offloading navigational computations onto it, can I say that the star is computing north, and does it continue to compute on a cloudy night? And just how big is a cognitive artifact? A sundial includes, mind it, the Earth and the Sun. Are these celestial bodies extensions of the mind when we read the hour of the day on a sundial? The fact that these questions are both legitimate and puzzling makes the extended mind narrative *hard to work with*. The narrative puts us in the position of raising questions that are legitimate and pressing, while at the same time, it tells us that no viable answers are forthcoming. As a consequence, it is difficult to make sense of it either explanatorily or (even more worrisome) descriptively. It is not clear which scientific use to make of the notion apart from the sense of awe that the image can inspire in us.

If, however, the narrative of cognitive extension is accepted, it instantly threatens to become a platitude such as we should merely decree that all cognitive artifacts are a part of our mind when we use them (and in many cases, when we do not use them as well). We will thus never be in the position to use the claim for making distinctions or formulating hypotheses as we would already know the answers ("What about this compass, abacus, or hard disk?" "Yep, just yet another part of my mind; *what else* could it be?").

Extended mind concerns are mirrored by concerns about the "extended body" view of artifacts (a view that has attracted less interest) and possibly supported by more stable intuitions. The blind person's cane is arguably an extension of her arm. But where can we stop the extending? Is a car an extension of my body? Is the cab where I sit now an extension of my body when it is driven by a taxi driver? Am I, the taxi driver, and the cab part of a "system," and does the system include roads and traffic lights? What about sitting in a subway car? And is the notion of a system psychologically relevant (or even useful)?

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

We can focus on a simple case of extended body such as prostheses, aids, and the highly quoted Cartesian cane of the blind. After a necessary training phase, prostheses are said to be "incorporated," as we feel the ground at the tip of the cane and not at the tip of our fingers. Even so, this means *not* that the cane has become *part* of the body but only that we have generated through training a complex and effective *representation* of the cane as part of the body. There is no urgent need to talk of an extended body of which the cane is a part. Likewise, our cunning ability to exploit environmental regularities does not require reconceptualizing the environment as an extension or as a part of the body. I am pretty skilled at using terrestrial gravity for walking and would fail when I use gravity on the Moon without training (or I could quickly readjust if I want to walk and not bounce). But what is the use of claiming that the *Earth's gravitational field* is thereby *part* of my *body*? (Note that there are so many other bodies influencing local gravity, such as the Moon and the Sun and beyond, even the supermassive black hole Sagittarius A* at the center of our galaxy.)

It looks as if the most sensible use of the theoretical ideology of the extended mind is heuristic: it helps us direct our attention to a class of phenomena that involves development and learning of sensorimotor skills inherent to body–artifact coupling. Thus, I'm happy with a *symptomatic* use of the extended mind ideology. But we definitely do not need to commit to the *ontology* that corresponds to this extended mind ideology, or, at least, not without a lot of circumspection and attention to the many epistemological consequences.

Thus, I see here what I shall call a *Distraction* (and a big one at that). In particular, the "offloading" idea blinds us to both the complexity of the interactions that occur between cognitive artifacts and the mind and of the internal reorganization of the mind/brain that makes it possible to use the artifacts under study.

> Methodological caveat: Distractions versus objections. Philosophers argue from definitions and examples, and criticize other philosophers' theories by alleging counterexamples that in turn hopefully lead to better definitions, further objections, rejoinders, in an iterative process. Much as I am happy with the methodology because it allows one to explore extremely detailed and ramified conceptual complexities, I feel the need to protect the possibility of preserving a train of thought from a potentially unlimited set of iterative objections. Thus, I'll employ the label "Distraction" when the objection it represents, though legitimate in principle, is relatively harmless to the core of the argument. (I actually enjoy Distractions.)

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

One further important worry is that the *phenomenology* of the extended mind is elusive. At a bare minimum, it is compatible with *feelings of reliance* on artifacts. In a sense, Plato had it all right when in *Phaedrus* 274b-275 he has the Egyptian god Thamus challenge the old god Theuth on the benefits of writing as an invention (Plato 2006):

> But when they came to letters, This, said Theuth, will make the Egyptians wiser and give them better memories; it is a specific both for the memory and for the wit. Thamus replied: O most ingenious Theuth, the parent or inventor of an art is not always the best judge of the utility or inutility of his own inventions to the users of them. And in this instance, you who are the father of letters, from a paternal love of your own children have been led to attribute to them a quality which they cannot have; for this discovery of yours will create forgetfulness in the learners' souls, because they will not use their memories; they will trust to the external written characters and not remember of themselves. The specific which you have discovered is an aid not to memory, but to reminiscence.

Stating "an aid not to memory, but to reminiscence," this passage can be read as a suggestion that writing is not expanding memory, that a biographical episode consigned to the written text is not of the same psychological kind as the mental memory of that episode. And on the other hand, the passage suggests that a *different* psychological item, reminiscence, is *enabled* by the written text. The specific *mechanisms* of reminiscence (the onset of *trust*) and the consequences of its systematic use instead of exercising memory (memory loss) are also outlined in the passage. The core computations of memory are deployed not for remembering an item but for remembering where information about that item is stored. The subject does *not remember* what she wrote in the notebook, and this is why she uses the notebook in the first place.

The extended, embodied, embedded, and enactive accounts of mental life are heralded as radical alternatives to the representational-computational account of the mind that this book endorses throughout. I think of them rather as challenges than alternatives. A sizable part of the theoretical effort undertaken here is to make sense of the challenges and to account for the data they provide in computational terms. This is what I propose to do when I discuss the *mechanisms* that involve cognition through extra-bodily objects. In the meanwhile, if you want to continue using an "E" word, I propose the expression "equipped mind."

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Ambiguities of Offloading; Displacing versus Genuine Offload

Offloading is our next critical target. Consider what happens when one uses a shopping list to do one's groceries. Why use the list? Because, by so doing, one does not have to keep in mind the various items on the list. Keeping them in mind has a number of costs: mental representations in working memory are volatile, there can be errors, and maintaining the representations prevents you from efficiently doing other things, such as engaging leisurely in conversation, planning your whereabouts, or singing a song. The corresponding cognitive advantage of the list is that it frees working memory, making it available for other tasks. If you have to keep in mind all the groceries, you are unable to do other intellectually engaging things, and you risk forgetting the groceries themselves.

But does this advantage support the strong notion of offloading, that is, the idea that as the memory task is offloaded to the list, the list is actually remembering things on your behalf? The answer is no. For a list to be used in the grocery store, some further engagement with the list is necessary (Goody 1978). One must for instance check off items as soon as one has bought them. Otherwise, one must mentally keep track of items deleted, and with that, the memory advantage vanishes, even though there is a likely asymmetry between having to keep track of unwritten items and having to keep track of invisible deletions of written items. Looking at lists as memory savers is reductive insofar as it prevents us from seeing other nice features they have, such as their *contractual* nature. They not only remind you of what you should buy but they try to prevent you from buying things you did not plan to buy. So, there is a rich cognitive life of lists. But as for maps, this life only unfolds through an interaction with the user.

Mind the offloading fallacy. From the fact that a cognitive artifact is *relieving you from certain tasks*, it does not follow that *it is performing those tasks for you*. Or at least, claiming that it does would require a lot (a lot) of arguing. Claiming so would mean claiming that lists remember, that roads know their way to the destination, that maps know *all* the roads to all destinations, and so on. On the other hand, the offloading metaphor is perfectly suitable for other cases, typically run in Mode 4 (M4). It can be safely claimed that pocket calculators do compute sums and products for you, or that your classmate does your assignment for you, or that your spouse remembers the shopping list for you. But we are not entitled to extrapolate

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

the structure of these (M4) cases to the structure of cases such as that of the list, the road, or the map. Cases of genuine offload (use of a pocket calculator) require a separate treatment as they do involve other (auxiliary) mental computations, those necessary to use them.

If we fend off the offloading fallacy, what are the options?

- A strong internal *displacing* claim would be that there is no offloading at all and that the cognitive load is internally displaced to another module or system *in the brain* thanks to the interaction with a cognitive artifact.
- A weak thesis is that *even if* offloading occurs, we *still need to explain internal displacement*, as it surely happens that some computations are displaced whenever a cognitive artifact (including a known offloader) is engaged cognitively.
- An attenuated scope of the notion is available, possibly to be replaced or complemented by a weaker notion, that of *relieving*.

Note that one can "offload" one's mind by *doing* the thing one needs to do in order to free working memory. I keep seeing a skewed painting, I keep telling myself that I should do something to it, and I obsess about it until the point comes when I just hang it straight. At this point, one can claim that I have freed my memory. Is this a case of offloading? In this case, no cognitive artifact is assisting my cognitive task. Rather, the cognitive task just vanishes as I did what I intended to do and what I kept repeating to myself I should do.

Representational Epiphanies and Natural Signs

Maps represent. Are they arbitrary? The tension between natural signs and arbitrary representations has attracted a lot of interest in the literature. Here I follow the approach by Dretske (1986; 1988), who distinguishes between indication (or information) and representation (MacFarlane 2015). Indication is independent of our purposes and needs (a compass needle indicates the direction of the magnetic North Pole no matter what). Indication depends on environmental circumstances (large metal bodies nearby modify the direction of the needle) and can be indeterminate (the movement of the needle can depend on a change in the Pole's position or on a change in the compass's position). Indication leaves no room for "mis-indication." Representation is, on the other hand, open to misrepresentation (a map

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

can represent Paris north of London) and is sensitive to description (a map can represent London, London is the capital of England, but the map may not represent London as the capital of England). As mentioned in the beginning of the chapter, we need the notion of representation above that of indication because it is the only door we have to formulate hypotheses on the computational workings of the mind.

The representational paradigm opens a number of avenues in the characterization and explanation of epistemic artifacts. An example is what I shall call a *representational epiphany*, which appears as a key mechanism for the genesis and design of artifacts, notably artifacts that build on other artifacts. Consider one hypothesis about the origin of drawings. According to (Kennedy 1975), drawings were discovered, not invented. For instance, when looking at the profile left by a leaf on a stone (figure 1.2), someone discovered that *contour lines are enough to induce recognition*.

Spotting faces in clouds (Gombrich 1960) is another example of a representational epiphany. Visual artists have long been interested in epiphanies, and attention to them is part of pictorial textbook recommendations: stains on walls, inkblots, the face in the Moon, animals in clouds. One example for

Figure 1.2
A representational epiphany. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

all is the story of Dibutade inventing painting from shadow tracing as told by Pliny the Elder (Pliny 1967). Epiphanies are different from the (over)interpretation of "magical" or "telling" signs by soothsayers, as they do not require complex interpretation or esoteric knowledge. They are just that, epiphanies!

The phenomenon of an epiphany is not limited to the visual modality. The human voice can imitate a massive number of different and strange sounds (Lemaitre et al. 2016). Some of the representational properties of the human voice are likely to have been discovered through a series of representational epiphanies.

Importantly, natural signs and representational epiphanies are orthogonal to each other. Natural signs such as shadows, silhouettes, and tracks may or may not generate representational epiphanies. For shadows, sometimes they do and sometimes they don't (Casati and Cavanagh 2019). A "recognition shadow" allows one to tell the identity of the shadow caster. But most shadows that one casts are silent about the caster's type or individual identity (as is the blobby shadow one casts on the ground at noon under the Sun). Conversely, representational epiphanies may be generated by items that are not natural signs. A random distribution of inkblots may be epiphanic in its sudden looking like a pet but is no natural sign of the pet. It is a natural sign of the event that generated the spread of the inkblots, but it is not epiphanic of that event.

Taking into account the point made previously, representational epiphanies are conscious mental episodes. The visuomotor system may unconsciously retrieve from the visual properties of a line traced on paper the gestures that lead to the drawing of that line. This retrieval can be measured in various ways (Pignocchi 2010), but we normally do not have conscious access to the tracing gesture. For instance, we have no clear idea of what the gesture looked like or in what sequence it was performed. Representational epiphanies may play a role in what Karmiloff-Smith (1992) has called *representational redescription*. When children learn to draw, they start tracing circular shapes that are later reinterpreted as a base for the representation of faces.

Finally, representational epiphanies are conscious, involuntary recycling of natural signs. They can be contrasted with *neural recycling*. A prime example of neural recycling is the colonization by the reading ability of cortical areas devoted to object and face recognition in the left lateral occipitotemporal sulcus (Dehaene and Cohen 2011). It is not by chance that alphabets in the world use a restricted number of graphic primitives, as these are the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

particular shape and shape features that the Visual Word Form Area was tuned to before the invention of alphabets.

We will see in chapter 2 in what sense maps are epiphanic.

Varieties of Cognitive Artifacts: Definitions and Classification versus Cognitive Mechanics

Maps are cognitive artifacts, and some proposals for classifying cognitive artifacts are available in the literature, from Hutchins (1999) to Heersmink (2013) to Carrara and Mingardo (2013), among others. I stand by a minimal classification without making any claim as to its completeness or adequacy, as taxonomic endeavors are of intrinsic but limited interest for our purposes. Artifacts in general are a mixed bag, and there is an abundance of counterexamples for each proposed characterization. We do not have stable intuitions about artifacts' nature and identity as the many proposed "solutions" to the Ship of Theseus puzzle continues to suggest (Casati 2011). *Artifact words* are a mixed bag too. They can be quasi causal ("aspirin"). They can be ruled by family resemblances such as "chair," where "an armchair is a chair" in English but not in some other languages. Thus, in Italian and French, an armchair is *not* a chair ("Una poltrona non è una sedia"; "Un fauteuil n'est pas une chaise"). Other *artifact words* follow internalist-criteral intuitions ("Knoll chair") (Marconi 2013). As a matter of fact, classifications of artifacts tend to be centered on the contingent interests of those who present the classification in order to highlight a certain artifact or a certain artifact function, thereby (mildly) distracting from the complexity of the cognitive mechanics.

Here, the minimal classification proposed distinguishes, within the overarching category of artifacts, the class of *cognitive* artifacts that intersects the class of cognitive *representational* artifacts. A cognitive artifact is an artifact that purposely weighs on the quality of our epistemic states (or of our metacognitive feelings). A representational artifact is a cognitive artifact that has representational content. Mirrors, periscopes, and eyeglasses are cognitive nonrepresentational artifacts (Casati 2012). Pictures, writings, diagrams, and maps are cognitive representational artifacts (figure 1.3). The classification of artifacts leaves room for representational noncognitive artifacts, such as the internal code of a computer that is not supposed to be accessed cognitively; in what follows, we will ignore this further class of artifacts.[3]

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

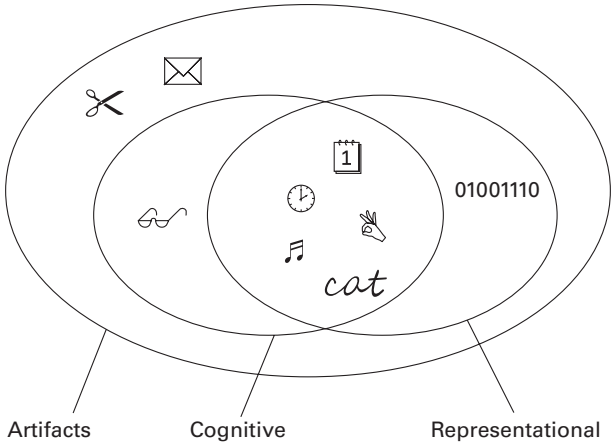

Figure 1.3
Varieties of artifacts. Image credit: RC.

Representational artifacts are characterized by a *vehicle* (the physical item in the world that does the representing: a piece of paper with colors and letters on it) and a *content* (what the vehicle represents or is about: a territory, a person, an event). Representational artifacts are further classifiable according to *representational systems*, that is, ways to organize information to produce a certain representation (Haugeland 1991; Kulvicki 2006). When it comes to maps, we can distinguish them, at least intuitively, as different from diagrams, pictures, and written text, while also noting some similarities. The boundaries between these categories may be porous, and part of the conversation on representational systems is about how to label artifacts, at least in certain limit cases: are maps a kind of picture? Are they akin to diagrams? These questions will keep us busy in chapter 3.

Cognition of Artifacts versus Cognition through Artifacts

We use cognition to interact with artifacts. For instance, we detect an intention behind the shape of a hammer and the stuff it is made of, which in turn helps us understand how to use it. This goes beyond affordance perception: the hammer head affords grasping. In fact, it is sometimes more comfortable to grasp the hammer by its head than by its handle, in particular if you want to carry it for a long time, but this is not how the hammer is meant to be used.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Intention detection and the cognition *of* artifacts, generally speaking, apply to artifacts across the board as conditions of their use. This claim also holds for cognitive artifacts such as maps. But maps, as well as written texts, are also tools for cognition; there is not only cognition of them but also cognition *through* them. Once you understand how a map works, you can use the map to understand a given territory.

Abstraction: Systematic or Ad Hoc Omission

As I write this book, I am sitting on the fourth floor of a building on Amsterdam Avenue in New York City, New York, United States. The place where I am is registered on the map I'm reading as one of the two tallest places in Manhattan. How is it that I see many buildings around me that are higher up in the sky? The elevation of a place is normally rendered by measuring the elevation of its footprint rather than its height. At the ground level, my location is the tallest in NYC, but not at the fourth floor. This distinction appears paradoxical or just confusing, as the fourth floor is way above ground level. By considering only the footprint, we *abstract* from what is built on it. The reason we do this abstraction is, say, that we need information about the watershed (figure 1.4). If, on the other hand, we were to fly drones into the city, the ground-level information would be insufficient, irrelevant, or confusing. In other words, abstraction in maps is defined by

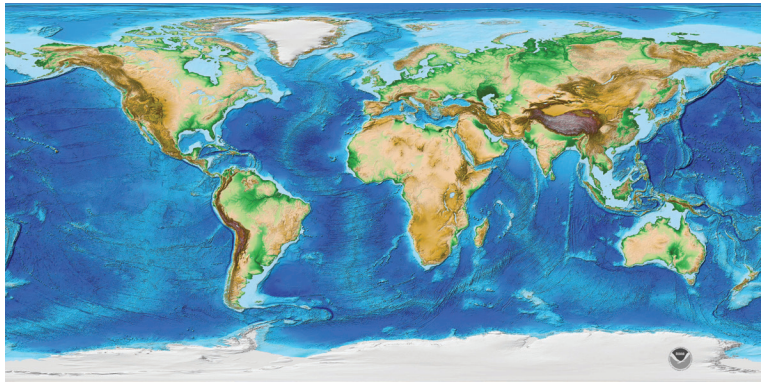

Figure 1.4
An elevation map. Image source: NOOA.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

not considering certain facts or bits of information when representing a territory: its omission from a representation.

Abstraction can be systematic or partial. Maritime maps only represent some salient features of a coastline. Sketched maps just focus on waypoints (Denis 2018). Monument maps of cities represent the footprints of most buildings and sporadic 3D renderings for important monuments. These are ad hoc omissions. Abstraction can be more systematic, as with the footprint map of elevations where everything above ground level is omitted, or with a roadmap where only highways are represented.

Mind an abstraction Distraction! When we abstractly represent a certain territory, do we represent an abstract territory? Do objects have abstract parts that are what get represented in abstract representations? Casati and Soldati (1995) argued thus, but one can aim for a less demanding line. Here abstraction is considered a property of representations. The bottom line of this consideration is derived from the fact that if one represents an object abstractly, it does not follow that one represents an abstract object.

Map and Territory

Representational artifacts have both a vehicle (for maps, the foldable paper sheet) and a content. I endorse a basic distinction between map and territory (cf. Searle 2018). Accepting the distinction is compatible with accepting strong and potentially controversial claims about the relation between maps and territory, including claims that maps create their respective territories or that territories are made from maps (some of these claims will be discussed in chapter 9). A limit case (discussed at the end of chapter 4) is that of Borgesian maps where maps can spatially coincide with the territories they map, but coincidence is not identity.

Mind the Distraction: "A map is its own map."

Can there be maps without a territory? A map purporting to represent Middle Earth is intuitively still a map in the same way that "The present king of France is bald" is intuitively still a sentence. Goodman (1968) distinguished between horse-pictures and pictures of a horse in order to preserve the ability to classify pictures. A horse-picture differs syntactically from a unicorn-picture; but if there are pictures of horses, there are no pictures of unicorns. This is because there are no unicorns, and no unicorn-picture is a picture of a unicorn. And some pictures of horses fail to be horse-pictures

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

when you are unable to tell what they purport to represent. We can reuse this distinction and the attached terminology to state that Middle Earth maps are not maps of Middle Earth.

The Mechanics

The remainder of this chapter presents some of the mechanisms that are deployed in our interaction with cognitive artifacts. I do not pretend to be exhaustive. Rather, I point out some examples that are interesting both per se and because they reveal what a working theory of the cognition involving artifacts can be, in particular as they guide our discussion of maps.

Displacement and the Schumann Principle

> In my opinion it would be a good plan to express these musical terms by a series of symbols, which have more affinity with the notes than have the letters. How much more quickly the eye takes in the symbol < than its Italian verbal equivalent *crescendo*; there is a pretty charm in the different signs, bow, lines, etc. used with musical characters, and I have oftentimes thought that the manner of illustrating a work by symbols of expressions throws more light upon a composer's aesthetic culture than do the very tones.
>
> —Schumann, 1983, 48–49; my underlining.

> Again, it seems to me a matter of no small importance to have ended the dispute about the Milky Way by making its nature manifest to the very senses as well as to the intellect.
>
> —Galilei, 2013

Both Galileo and Schumann adumbrate what appears to be one of the most powerful building blocks of our interaction with cognitive artifacts: the *displacement* of certain tasks from M2-like processes to M1-like routines (figure 1.5).

We can see displacement at work in a simple example. Suppose you are looking for the trains leaving between 9:00 and 11:00. You are presented with four timetables, each displaying the same set of facts represented by the times trains leave. But they are shown in different presentations: (Bag), (Order), (Highlight), and (Chunk). From these presentations, you shall notice (table 1.1) a difference in retrieving the desired set of times.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

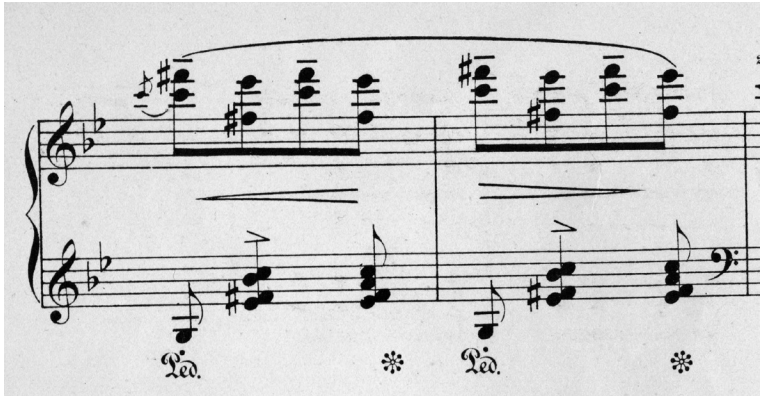

Figure 1.5
A musical crescendo, followed by a diminuendo. E. Grieg, Berceuse, Editions Peters, end of nineteenth century. Author's collection. Image credit: RC.

Table 1.1
Which Trains Leave Between 9:00 and 11:00?

(Bag):	10:20, 8:30, 11:40, 13:30, 9:25, 10:35, 16:40, 8:15
(Order):	8:15, 8:30, 9:25, 10:20, 10:35, 11:40, 13:30, 16:40
(Highlight):	**10:20**, 8:30, 11:40, 13:30, **9:25**, **10:35**, 16:40, 8:15
(Chunk):	8:15, 8:30, 9:25, 10:20, 10:35, 11:40, 13:30, 16:40

Larkin and Simon (1987) would claim that the four timetables are *informationally equivalent* but *computationally nonequivalent*. They all say exactly the same thing—the times trains leave—but the costs of extracting information differ greatly among them. On a different but compatible line of reasoning, Tversky et al. (2007, 56) would distinguish between the principle of congruence ("The structure and content of a visualization should correspond to the structure and content of the desired mental representation") and the principle of apprehension ("The structure and content of a visualization should be readily and accurately perceived and comprehended"). Now, these overarching distinctions do not provide us with the specific mechanics of displacement. What are the computational differences that matter for apprehension? Why does it take longer to search the presentation (Bag) than in any other style? What advantages are available here?

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

If you think about your behavior when faced with (Bag), unassisted by paper and pencil, you'll notice that you have to:

- pay attention sequentially to each item in the series; then
- explicitly mentally label it as a positive or a negative; then
- store positives in memory; and possibly
- revisit the series for double-checking.

That means that both targets and distractors are time and again processed and reprocessed. You must remember what the target's content is, or where an encountered target was, in case you forget the content. Both spatial and short-term memory are actively engaged and put under pressure. You pay attention to each item. All this indicates that your cognitive activity is located in an M2-like area. In the other three formats, powerful and fast computational routines of the visual system take up part of this chore, shifting it to M1 modes of operation. The visual system loves ordered series, color differences, and chunks, and can instantly pinpoint the true positives as they are packaged in a visual-system-friendly format (Healey and Enns 2012). Displacement of cognitive tasks is one of the key mechanisms of cognitive artifacts. The implicit design of (Order), (Highlight), and (Chunk) is an orchestration of M1-operating routines that are more effective than the M2 work done when searching in (Bag).

Displacement is interesting in two other ways. First, displacement does not require learning, and thus is extremely effective in giving designers of cognitive artifacts great leverage. Second, displacement keeps computations within the boundaries of the brain, moving them from one part of the brain to another; the design of the artifact triggers the new computations and makes the old ones moot.

The nonnecessity of learning distinguishes displacement from other shifts of cognitive load within the brain. For instance, by gaining expertise about paintings, art dealers train their brain so that certain complex step-by-step (M2) judgments about the quality of a painting are automatized and get compressed into intuitive, perception-like assessments (M1). Art dealers involve different parts of the brain when they examine paintings (Kirk et al. 2009), but the shift to automation in their observations requires a fair deal of learning.

One reason to highlight or underline items on a list (and thereby shift their processing to the visual system) is that when you read the list *without*

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

highlights you actually read it and likely get distracted from the task you had in mind when you began reading. This task may be a simple search task or it may be a larger task that includes a search task.

Shifting mechanisms are used to deal with mathematical formalism, that is, equations and logical formulae. In learning to work with equations and logical formulae, we take advantage of visual shortcuts that with time become second nature (e.g., you can "see" an acceleration in a velocity diagram when you see a certain change in slope or you can see that a certain factor in an equation can be replaced by a complex formula). It is however important that a complex learning process takes place here. In logic, *normal forms* are a way to reorganize information and content so that it is computationally fit for a certain task. For instance, an argument is in normal form when its conclusion is placed after its premises. Quantified formulas are in prenex normal form when all the quantifiers are gathered to the left of the formula. Normal forms are task dependent, there is no such thing as "the" normal form for a given item. Returning to our example above, a timetable is arguably in normal form when the times are in (Order) mode, not in (Bag) mode. At this point, one can discuss whether descending order is better than ascending order. The choice will amount to setting a new task-relative normal form. We'll see in chapter 2 how to apply this notion specifically to maps.

From Brakes to Barriers

This preliminary characterization may suggest that the main job of cognitive designers is to make cognition more fluid, move it from cranky M2 to smooth M1. However, the opposite is sometimes desirable. We may want to prevent cognizers from jumping to conclusions. Thus, we try to provide cognitive *brakes to* force them to abandon an M1 routine and think twice about a certain problem. Shifting is not just from slow, painful M2 to quick, fluent M1. Often, we ask students to use paper and pencil to analyze problems whose intuitive answer (or "the first thing that comes to mind") is under the spell of powerful cognitive biases. "The ball and the bat together cost 110 units, the bat costs 100 units more than the ball, what is the cost of the ball?" In order to inhibit the powerful, intuitive M1 (wrong) response "10," we use paper and pen and knowledge of rules, that we apply in M2 style: ((bat + ball) – (bat – ball)) / 2. Both logic's formalism and arithmetic representations act as cognitive brakes.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

We return to the train timetable. If (Order) provides cognitive advantages for the quick search task, (Bag) slows down the search task. Relative to a search task, (Bag) works like a *Barrier*. There are many situations in which one needs barriers. Designers of ATMs recognized soon enough that closing the operation cycle of cash withdrawal with withdrawal of the debit card generated unacceptable consequences—people walked away with the money and left the card in the ATM. Thus, designers ended up inverting the order of withdrawal: "Remove your card in order to obtain the cash."

Barriers can be involuntary. Button panels on elevators are maps when in order (and are oriented according to the position of the floor), but at times they are rearranged in (Block) mode and act as barriers (see chapter 8). Barriers such as warning signs block fluid, but potentially counterproductive routines that capture the user's attention and enforce a shift to M2 controlled procedures. Camouflage is an example of a strong barrier that takes advantage of biases in the perceptual system to group together items that are close and alike. Once camouflage is in place, it blocks the perception of the camouflaged object (Sperber and Hirschfeld 2004).

Shunting

Another extremely powerful mechanism of cognitive artifacts is *shunting*. When looking at a line drawing representing a chair, early phases of information processing are *shunted*: line extraction is directly interfaced with the recognition of a chair (figure 1.6).

An enormous cognitive advantage (quick recognition) is provided by the very parsimonious use of simple and uniform lines. The strong explanatory hypothesis posited by Cavanagh (1999) is that at some level of visual architecture, information is stored in a format that only specifies occlusion boundaries—precisely as line drawings do. This format is the mental target of line drawings. Thus, given the hypothesized architecture, drawings work efficiently because they allow shunting of information. This fact teams up with the simplicity of the production of drawings, as opposed to other representational styles such as photorealistic images and silhouettes (see Willats 1997): drawings are both representationally effective and easy to produce.

Another example of categorization-related shunting is differential coloring of parts of artificial objects. Very few animals have limbs whose color dramatically differs from that of their main body. A chair with legs and seat

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

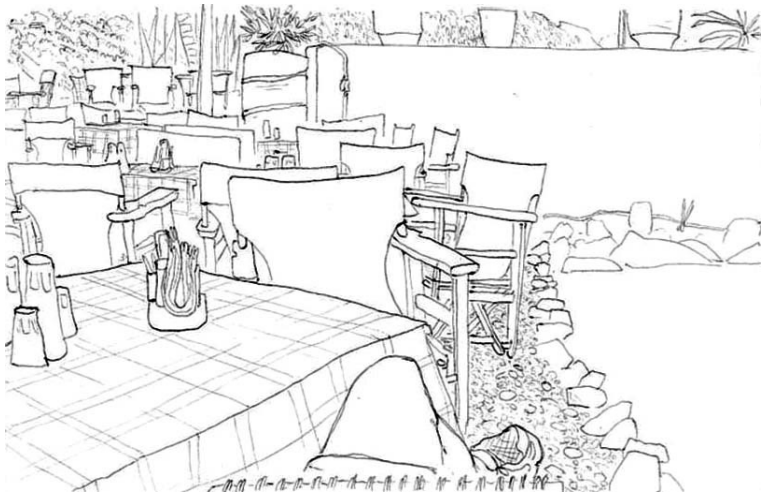

Figure 1.6
Chairs. Image credit: RC, 2013.

of different colors is easily recognized as an artifact. Structured coloring here is the cognitive artifact that enables faster recognition of a noncognitive artifact.

Bridging

External representations can act as *bridges* between representational systems that have no biological or evolutionary reason to team up. Bridging is the key, for instance, to understanding how reading and writing work. Letters written on paper are the external bridges between the visual recognition module and the phonetic/articulatory module. These two modules had neither evolutionary ground nor evolutionary history for working together (Dehaene 2009). The modules team up for those who can read after years of tireless, and often painful, education and operant conditioning establishes the relevant neural and functional connections between them. To team up, the modules require written letters and the relevant external bridges (figure 1.7). And when they do team up effectively, you are not able to simply look at the patterns on this page and see mere scribbles or decorations. In fact, you cannot help but read what is written here. Thus, your mind has perceptibly moved from an M2 slow deciphering to an M1-type fluency.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

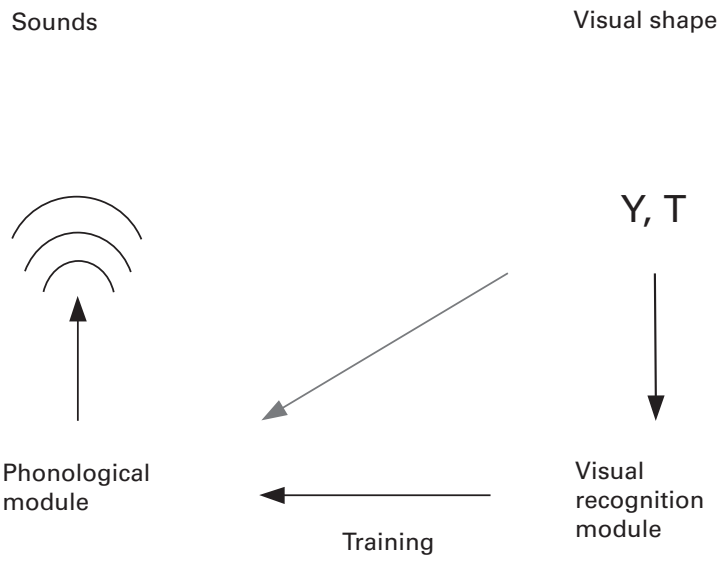

Figure 1.7
Writing is an external bridge between visual recognition and linguistic sound production. Image credit: RC.

Compare this process to what happens when you try to read in a different writing system you only have begun to learn. For those who do not read Greek, "Αθήνα" requires M2 action: look at a list of alphabetic symbols and spell the word letter by letter.

Floors and Ceilings

Kulvicki (2010) systematized two important syntactic maneuvers one can perform on certain types of representations. Consider the temperatures represented in figure 1.8.

The same geographic territory is represented three times, with temperatures in Celsius in different quadrants. In the top left part of the image, the temperatures are expressed by figures. In the top right, the designer "raised the abstract ceiling" by coloring the figures according to a grayscale-temperature scale, thereby making abstract information more easily recoverable from the representation. Through the top right numeral coloring, we can "see" that there is a temperature gradient running diagonally in the represented territory without having to draw M2 inferences as we would

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

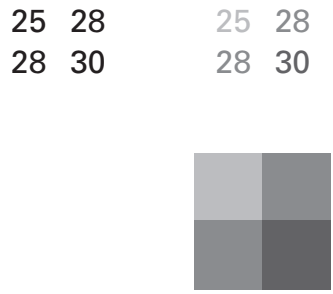

Figure 1.8
Temperatures in a parcel of territory and two syntactic operations: raising the ceiling (top left to top right), then the floor (top right to bottom right). Image credit: RC.

if there are only black digits (i.e., by making comparisons between digits: "28 is to the right of 25 and 30 is below 28, therefore there is a temperature gradient"). At the bottom of the image, we took the grayscaled digits representation and raised its determinate floor. While we no longer access the exact values of the temperatures, we do speed up the assessment of the temperature gradient because the whole computation is now carried by the visual system. Though we lose some specific information about the absolute values, we immediately access the relative values.

The Mechanics Impose Design Constraints on Cognitive Artifacts

We have enlisted a number of tools that allow us to describe and characterize some aspects of the life of cognitive artifacts. In particular, we can now see that cognitive mechanics impose huge constraints on the design of cognitive artifacts, that is, of the way manipulating *vehicles* (e.g., *raising ceilings* or raising *floors*) makes it possible for us to understand and intervene on contents, to move certain mental operations from an M2 mode to an M1 mode, or vice-versa. Constraints also induce *representational epiphanies*, *bridge* representational modules, and *displace* or redistribute cognitive load across artifact architectures. In general, M3 intervenes in M2-like ways by orchestrating different M1-like activities, and in M1-like ways through certain fine-grained interactions with artifacts.

These general points have specific applications to maps and our interactions with them. When you use a map for navigation, you use an M2-like process to establish a complex back-and-forth between map and terrain.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Once a relevant bit of information has been gathered, you make progress, update your position, and then you do not go back to look for that piece of information again. Moving your head back and forth from map to terrain to gather information with your internal spatial mapping system, on the other hand, is performed in an M1-like way. To create a good, usable map is to be able to keep in balance two sets of constraints (Tyner 2014): the map must be accurate, relative to a standard of accuracy, and it must be readable and interpretable by people with certain cognitive abilities. How the map is capable of representing its territory and how it can be read and used is what the toolbox we assembled will allow us to assess in the rest of the book.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

2 The Central Features of Maps: Systematicity, Hypergenerativity

A Theory of Maps

Maps—once more, we mean here "external" maps, things outside the body that you can consult on a screen or unfold on a chart table and show to other people—are central cultural items in human culture and history. Every day, billions of maps are used to organize knowledge about the territory around us and to plan and execute navigation practices. Unsurprisingly, a sizable amount of literature has been devoted to maps.[1] And yet their mechanics is not fully understood, and there does not exist a consensus about what items can qualify as maps. In this chapter, I hone in on some tools specific to understanding maps while also showing how the mechanics of cognition both *of* and *through* maps contributes to an understanding of some key aspects of mental architecture. The next chapters will build on the theory presented here, pursuing both a descriptive project (how maps work) and a normative project (what a map is or is not, and what is and what is not a map).

Maps, as they are understood in this book, are representations where you can physically *locate on them* the *representation* of something's physical location in a *territory*, *in a certain systematic way*, and they are such that they have the *same dimensionality* of the territory they represent. (This applies to spatial locations as well as to temporal locations, which will be further discussed later.) Systematicity is a first interesting element because it is an essential part of what makes maps maps as opposed to non-maps. Other elements such as same-dimensionality—admittedly, a nonstandard theoretical imposition amounting to the claim that maps only "see" certain superficial features of a territory—require an independent discussion, whereas

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

representational location concerns the usability of maps and is discussed at a later point in chapter 4.

Maps are tightly organized spatially. We intuitively know that systematicity is inherited from the rich spatial features of the *vehicle*. Consider, for example, a paper map: if you see on it that an area labeled "Bourgogne" is part of an area labeled "France," you are entitled to infer that Bourgogne is part of France. If you see that the area labeled "Belgium" is larger than the area labeled "Luxembourg," you are entitled to infer (prima facie) that Belgium is larger than Luxembourg. If the area labeled "Belgium" is in contact with the area labeled "Luxembourg," you are entitled to infer that Belgium shares a boundary with Luxembourg. And so on and so forth.

The theoretical framing of the systematicity of maps is provided by mereotopology (Casati and Varzi 1999; Cotnoir and Varzi 2021), that is, an integration of the theory of parthood relation (mereology) with the theory of wholeness or spatial integrity (topology). These theoretical tools are best understood in the light of what they can "see." A general theory of objects will tell you which things are, and which things are not, objects. Maybe your theory of objects tells you that an object is the referent of a common noun (e.g., "mug," "handle"). So, your theory of objects tells you that one thing is a mug and *another* thing is that very same mug's handle because you have a name for the mug and you have a name for the handle. (Other theories of objecthood exist, but here we use this simple example to make the point that follows.) Can we see more when we look at a mug and its handle? Mereology, from the Greek "meros/μερος," meaning part, is a theory of parthood relations. If you do not have a theory of parthood relations, you can see the distinction between a mug and the very same mug's handle, but you cannot see the *relation* between the two, that is, the fact that the handle is part of the mug. Topology, from the Greek "topos/τόπος," place or location, is a theory of the particular organization of the space objects are in or occupy. If you do not have a topology, you cannot see the difference between the configuration in which the handle is attached to the mug, and a configuration in which the handle is not attached to the mug: a whole mug and a broken mug are mereologically indistinguishable. In *both* configurations, the handle is part of the mug, so mereology cannot see the difference between the two configurations. Topology is a theory of *wholes*, of what makes a collection of parts a whole object.

(Many splendid Distractions keep philosophers busy here: Is a mug without its handle the same as the mug with it? Is a broken mug still a mug?)

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Further theories, for instance, a theory of causal connection, will allow you to see the difference—in itself invisible to a theory of objecthood, to mereology, and to topology—between a case in which the handle just touches the mug's body, without being glued to it, and a case in which the handle is solidly connected to the mug's body (so that if you pull the handle, the whole mug follows.) Yet other theories allow you to see size differences that in themselves are invisible or only partly visible to any of the above: Burgundy is part of France and is smaller than it, so maybe you know about ordinal differences, but no information is given about how much bigger France is than Burgundy, as opposed to knowing that the former is 528,573 km^2 wide and the latter is 47,784 km^2 wide. There are many open issues about the exact relationship between all these theories (Is mereology just a subtheory of topology?), but for the present project, we would like to activate the distinction between mereology and topology so as to be able to see the difference between certain pairs of geographic entities: Alaska is part of the United States, but it is topologically disconnected from the contiguous United States—mereology cannot see this relation, but topology can; metropolitan France is disconnected from the islands and overseas territories that are part of France; Malta is composed of some disconnected islands; San Marino is in one single piece; there is one big connected ocean with many parts variously labeled Pacific Ocean, Atlantic Ocean, Bering Sea, and so on; and there are some seas such as the Caspian Sea that are disconnected from the big one-ocean.

Systematicity means that maps respect the mereotopological constraint. (Dis)connected areas of a territory are represented by (dis)connected areas on a map. Areas related by parthood relations in a territory are represented by areas on the map that are organized accordingly.

The inferential entitlements are highly modulated by various, at times distracting, factors,

(projection methods may generate artifacts: on Mercator maps, the area labeled "Greenland" is larger than the area labeled "Latin America," and legends warn that the area of Latin America is actually seven times that of Greenland; flat maps may not show the maritime boundary between Alaska and Russia as a boundary; Europe can be represented twice over on some maps, as shown in figure 2.1; and so on)

but these distracting factors are in general kept under control in the deployment of interpretive practices. The key element holds: maps are representations that essentially use spatial features of the vehicle to systematically represent spatial features of the world.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

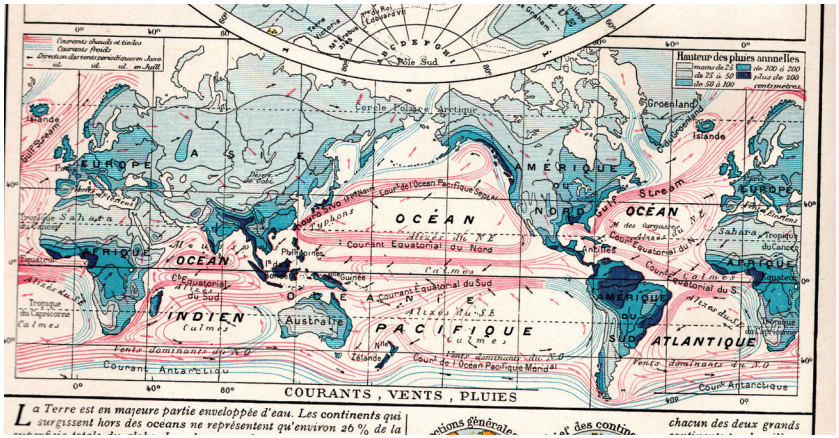

Figure 2.1
The territory does not contain two Europes. Interpretive practice blocks literal reading of the map. Image source: *Atlas classique*. Paris: Hachette, 1931. Author's collection.

Distraction: What if we had to cut a map into two parts? Some NYC metro maps present the north part of Manhattan on the left pane of a map and the south part on the right pane. If you did not mind this mapmaking choice, you would have to claim that topology is not preserved, as it looks as if the East River is now west of the Hudson River in a part of the map and east of the Hudson River in another part. Solutions to the Distraction: here we have *two* maps, one for each part of Manhattan. Countless other Distractions still lurk in this discussion. The East River is not actually a river but a strait. What if Bourgogne conquers France? Why is north up and south down in most maps? And so on. Distractions they are! More to come in what follows.

We now understand that maps are spatial vehicles that systematically use space to represent spatial content. Space as it is represented in maps is mathematically so well behaved that one may be tempted to climb up the abstraction ladder. Maps, one hastens to say, are just *mappings*! But we should resist this.

Maps versus Mappings

In mathematics, a mapping is a function, that is, an abstract device that takes every element of a given set as an input and then outputs a unique value from the elements of a given set. At a well-set dinner table, there is a mapping from each and every fork on the table to exactly one knife on the table. It should be noted that mappings are not the same as maps. Some

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

mappings are not maps in the sense that interests us. There is a mapping between my age expressed in years and my age expressed in lunar cycles. There is a mapping between the number of persons in a room and the number of nostrils in that room. There is a mapping between the side of a square and its area. None of these is a map. But are maps mappings? In a very abstract and innocent sense, they are. Maps are representations, and not only representations can be described as functions that take representational vehicles as inputs and output representational contents; in the specific case of maps, each and every area on the map corresponds to a unique area in the territory. But there is no particular use for this very general notion of mapping, and issues around it may just be terminological.

Two interesting cases: (1) a data set that generates a map (e.g., in a navigator) is a mapping, not a map, of the territory it represents [although the physical way data are arranged in the set (e.g., in tabular fashion) may authorize the claim that some data sets are maps], and (2) *place cells* in the hippocampus appear to provide a mapping of visited places without being organized in a map.

If we bar the unwelcome equation of mappings with maps, the way is open to an account that is bound to make potentially interesting predictions. An interesting prediction is a prediction that is relatively counterintuitive. Predictions that interest us concern contrast classes. This means that some items not normally considered maps will be recategorized as maps, and conversely, items considered "maps" are no longer maps (chapter 3).

Maps, Mappings, and Projections

If the technical notion of *mapping* is too large, the notion of *projection* is too narrow. Historically, mapmaking has been aware of the fact that a curved surface cannot be flattened without costs (stretching or cuts) and that some systematicity could be imposed on an otherwise arbitrary choice of where to place the graphical costs (Should we oversize Russia? Cut the map in the middle of the Pacific Ocean?) by using a projection, meant to be an optical projection from a viewpoint or light source onto a surface. Consider Mercator: a light source is placed at the center of a globe representing the Earth, and the shadows of all the points on the globe are cast on a cylindrical surface that wraps the globe and is tangent to its equator.

This optical interpretation does not work in the case of the azimuthal equidistant projection (used in chapter 6; as an example, consider the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

emblem of the United Nations). There is no "shadow" or "viewpoint" projection method that can take the globe to a flat plane where distances are preserved from a chosen point. What you do is consider the globe as covered with spaghetti connecting the chosen point to its antipodal point (e.g., the North and South poles, when the spaghetti are the meridians), then cancel the antipodal point and let the spaghetti drop radially on a surface (consider: spaghetti here do not have a thickness, so there are no gaps between them; the antipodal point degenerates into a line). This operation generates a mapping that preserves the mereotopology constraint (whereby it constitutes a map) but is not projective in the narrow sense of optical (or optic-like) lines connecting the source and the target of the projection (figure 2.2.) Therefore, either the notion of projection collapses on that of mapping or—in its optical/shadow reading—it cannot be applied to some types of maps such as those employing azimuthal equidistant projection.[2]

Maps and Analog Representations

Continuous variation in the territory corresponds to continuous variation on the map. Maps are considered analog, as opposed to digital (discretized), representations. Now the notion of analog representation is not completely

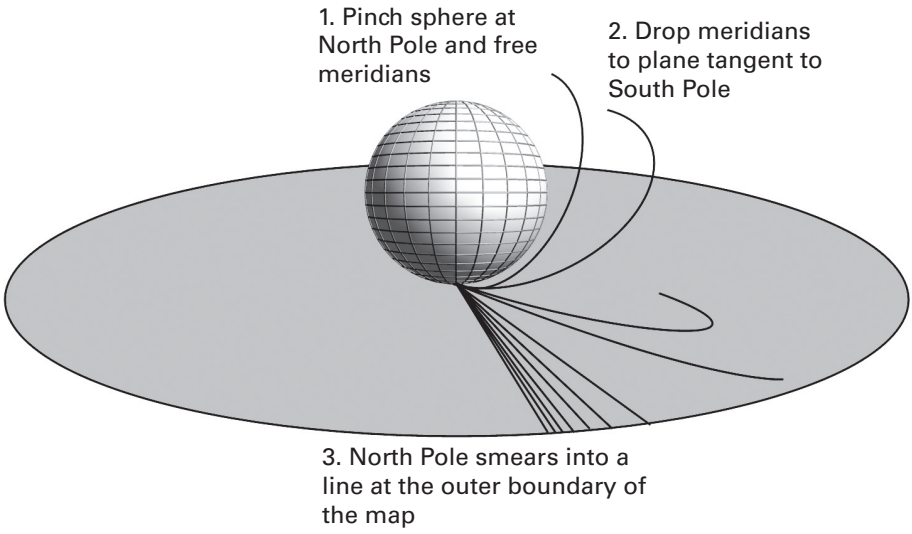
1. Pinch sphere at North Pole and free meridians
2. Drop meridians to plane tangent to South Pole
3. North Pole smears into a line at the outer boundary of the map

Figure 2.2
Azimuthal equidistant projection construction. Image Credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

stable or consensual (see Kulvicki [2015]; Dretske [1986]; Goodman [1968]; Haugeland [1981]; Lewis [1971]), and I do not plan to add anything substantive to the discussion at this moment. The simple working notion of maps makes maps a prima facie example of analog representations because of their sensitivity to certain minimal spatial differences. More generally, if analog representations are sensitive to minimal content differences, maps are just an instance of this, their specificity being that their content is the nested structure of spatial regions. Other analog representations track minimal quantitative differences that are either nonspatial (e.g., the second-by-second aggregated change in body mass of the population of Norway and the slant of wind pennants as a function of wind intensity) or, if they are spatial, are not concerned with the nested structure of spatial regions (e.g., the orientation of wind pennants as a function of wind direction).

A terminological Distraction: the capture of the digital/analog distinction by the digital-electronic/paper distinction. Technical Distractions: paper maps can be pixelized; nature may be discrete and not continuous.

Later on, we will address some finer distinctions between maps and other types of analog representations.

Methodology: Constraints on Map Dimensionality

A feature of maps that we take as essential in this discussion is that *maps must have the same dimensionality of the territory they map*. You have a 1D map of a 1D territory, a 2D map of a 2D territory, and a 3D map of a 3D territory. An example of a 3D map of a 3D territory would be a geological map printed in 3D with colors for different types of rock, or an anatomical model (figure 2.3).

Distractions pile here: What about flat[3] maps of mountainous terrains? (They map a surface that unfolds in 3D, but that is a 2D entity.) What about globes? (They are 2D maps that unfold in 3D space, representing 2D surfaces that unfold in 3D space.[4])

Whatever the interest of these cases, and of other limit cases one can find in the history of mapmaking or even ones we make up, we *still have to address how the constrained notion of maps functions*.

We call some of the 3D maps representing volumes of space "models," and I have no quarrel with that. But "models" fall in the same category as maps, and the only difference is the 3D parameter. So, for the present purposes, they are maps.[5]

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

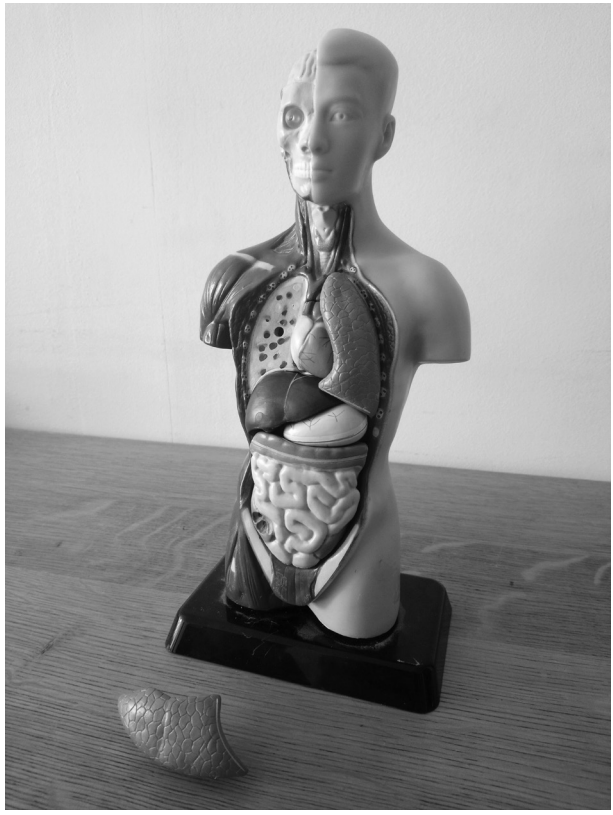

Figure 2.3
An anatomical model is a 3D map of a 3D territory (the human body)—at least when it is not disassembled. Author's collection. Image credit: RC.

A strong claim here would be that this constrained notion of a map in turn constrains what a mappable territory is. Not all types of territories can be mapped, nor can they be mapped by any type of map. There are limits to mapping, or at least to efficient mapping. A weaker claim with cognitive import is that maps *can only see a territory as having the same dimension* as the map. As mentioned in chapter 1, most city maps only see the footprints of buildings. They do not see a building's roof, floor, and so on. Physical maps only see the territory as its own projection on a plane.

Finally, same-dimensionality is not ontological homogeneity: music notation, as we shall see,[6] is a spatial map of a temporal territory.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

A Skeletal Semantics for Maps

The idea of systematicity can be made more precise. Casati and Varzi (1999) produced a skeleton of a semantics of maps that I utilize to make the point I want to underline on systematicity. The account of map semantics produced in 1999 has since been variously criticized in its general spirit and for some of its details (Rescorla 2009; Kulvicki 2015; Camp 2007; 2018; Aguilera 2016; Blumson 2014; Bronner 2015). For an assessment and some rejoinder, see Casati (2017); in this book, I focus on the basic ideas.

First, we proposed to work with an *idealized* notion of a map, one that allows us to highlight the *specific contribution of space to semantic systematicity*. The idealization practice is borrowed from formal logic and formal semantics, where certain fragments of language are abstracted as the shapes of items of a chosen vehicle. This abstraction allows one to elicit their contribution to the mechanics of inferences by various syntactic manipulations. In propositional logic, for instance, a piece of reasoning such as "Abi runs and Belle talks, therefore Belle talks" is rendered by A & B ⊢ B. Here, capital letters translate sentences, and logical symbols translate connectors between sentences, whereby we see that the important (necessary and sufficient) contribution to the validity of the argument is at the level of whole sentences and their connections. We do not need to care who Abi is or what Belle does; the argument is valid no matter what happens "inside" the sentences. This methodological move can be applied to maps in a way that respects and makes salient the contribution of their spatial structure to interpretive practices. This means that we do not study ordinary maps directly but rather their *translation* into *formal* maps.

Second, the *formal* map only keeps spatial regions as syntactic elements of maps, and we are open to any syntactic way to differentiate their content (to fix ideas, we use to that purpose an unlimited number of color hues). This entails a certain amount of *ingenuity in translating* ordinary maps, such as a map of the Upper West Side with all its symbols, into formal maps. For instance, we erase the label "Columbia University." If we want the map to represent a certain area in the world occupied by Columbia University, we need to color a certain map area in, let's say, brown. Then we would need to add a legend specifying that brown represents the property *of being* Columbia University. There are many (many) difficulties with coming up with an appropriate translation for most properties, which I detail at the end of the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

chapter. (As an example of the difficulty for an appropriate translation, in some cases, it may be that what the color brown refers to is not the property of *being Columbia University* but rather the property of *being an area occupied by a building belonging to Columbia University* or *being an area that is owned by Columbia University*, and so on.)

I do not claim here that I can address all concerns, but I will limit myself to asserting that limits to translation are limits to the expressiveness of formal maps—in the same sense in which limits to translation are limits to the expressiveness of, say, the language of propositional logic. Nobody "speaks" in the language of propositional logic, and we do not expect anybody to teach geography or to find their way by using a formal map (please don't). The important issue here is not whether formal maps have the same expressive power as ordinary maps but whether the expressive power they have allows us to use them to explain key interpretive practices related to ordinary maps, and in particular those practices that take advantage of the systematicity of the spatial features of maps.

> Distraction: Colors are visually loaded and cannot have the expressive power that semantics need. If we wanted to represent a human population and a bear population, we would have to overlap colors in certain areas of New Hampshire. What happens to overlapping colors? Human occupancy is assigned the color yellow, and bear occupancy is assigned the color red. Overlap would be another color, perhaps orange. Shouldn't orange correspond to yet another predicate, but wouldn't this exclude that area from being bear-occupied and human-occupied? This only intimates that there are technological limits to the color syntax. Colors, in this discussion, are just intuitively convenient placeholders for spatial occupants. The Distraction is dispelled by considering the layered structure of maps; evaluations are made on each layer where the problem does not arise. In this case, the orange on our map is a visual artifact. The overlapping area is not actually orange; it is instead both red and yellow.

I hesitate to consider as a Distraction the rich syntactic variety of symbols one can find on maps (such as lines and icons). However, for the limited scope I am setting to this chapter, they must remain a Distraction. Thus, I have little to say about accounts that focus on symbols other than spatial regions (MacEachren 1995; Pratt 1993; Camp 2007). The bottom line is that no matter how rich and varied the supply of symbols found on maps, we must *in any case*—and this is the crucial point—give an account of the semantics of skeletal, minimal maps whose only syntactic items are regions and colors. We cannot have a map if we only have symbols such as icons for campsites and printed names of cities, without at the same time providing

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

spatial regions on which to distribute those symbols. A list of ordered pairs in the form "<symbol, geographic coordinates>" is not a map. The following list exemplifying this form is therefore not a map.

<♥, 40°N,25°E>

<♣, 41°N,25°E>

<♥, 40°N,26°E>

<♥, 41°N,26°E>

While the following is a map of the same territory, by our lights:

♣♥

♥♥

The specificity of maps is that they represent space by space in a systematic way, and this should be the target of any theory of maps. In any case, it is the target of the theory defended here. Non-color symbols on maps (shapes, pictograms, and words) are themselves spatially structured, and their spatial structure (shape and size) *interferes* with map interpretation. In particular, the spatial structure of non-color symbols interferes differently at different scales with map interpretation. The figure for a road will use a different length-to-width ratio at different scales. The icon for a campsite has to be placed on the map, but where exactly is the campsite? Names of cities follow a different placing convention (e.g., the city is immediately to the left of its name's starting letter, unless the name interferes with other features on the map) than the names of countries (e.g., the country name is somewhat centered in the area denoting the country). All of these interferences, a form of spatial stroop effect, reveal the underlying power embedded in the spatial structure of the map.

Evaluation Atoms

The next step in the build-up of a semantics consists in finding an adequate "atom" for evaluation. In propositional logic, this is the atomic sentence ("A", in the example above, representing a natural language sentence like "Abi runs"). Once we assign it a truth value, sentences that are composed of atomic sentences are assigned a truth value by virtue of some rules of compositionality that transfer truth up the complexity ladder. This is based on the configuration through which molecular sentences are assembled. In

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

formalizing map semantics, we decided that a good point to start is, intuitively, what you get if you consider a map as *printed one color at a time*, each time with a color on a different sheet. Suppose you take the color black to mean *is Swiss territory*; if you took a pass at just black in the printer, you would see just the shape of Switzerland on a white background (figure 2.4). This would be a *map layer* (or *stage*, as we call it). The layer would be true whenever both the area in the world corresponding to the black area on the map layer is indeed Swiss territory and the area in the world corresponding to the white (non-black) area on the map layer *is not* Swiss territory.

Here there is a complication related to this "absence intuition" that we can consider briefly for our present purposes (see Bronner [2015]; Rescorla [2009]; Kulvicki [2015]).

We originally claimed that map colors work like predicates. When spread over an area, they generate something like a sentence. If a certain area is colored black, then the corresponding part of the territory is said to have the property expressed by the color black. Here, we endorsed an *absence intuition*. The part of the territory corresponding to the area that is not colored black is said not to have the property that black expresses. Predication in language does not work like this, however. If you say, "Fido sleeps," you are not thereby stating that other dogs do not sleep. But the differences between maps and language are not so radical. In restricted contexts, for example, introduced by questions such as "What do our dogs do?" the response "Fido sleeps" can be taken to entail that the other dogs are not asleep. Maps are best equated to this restricted context as giving very limited responses to very restrictive questions.[7] Their best analoga in natural language are

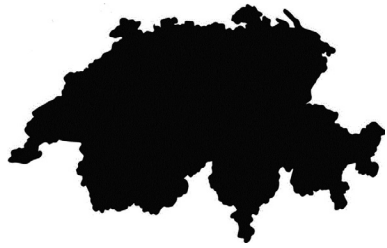

Figure 2.4
A map layer representing just Swiss territory and neighboring non-Swiss territory. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

probably *inventory lists* of items that respond to very restrictive questions in restricted contexts ("What are the books on the third shelf?").[8]

Historically, maps and linear textual lists of items are not so remote from each other. An itinerary is a way to solve a navigation problem by compressing it to a 1D line; a list is a 1D representation. Ancient navigation predominantly followed coastal routes. They were represented by peripli, ordered lists of coastal towns, with distances between them. Land travel was represented by itineraria. Some important historical maps, such as the Tabula Peutingeriana (a thirteenth-century copy and compilation of ancient Roman items), found it convenient to preserve the linear information encoded in itineraria and peripli, to the detriment of metric information about geographic regions.

To conclude our semantic architecture, a molecular map is the superposition/logical conjunction of these various map layers (one with black for *is Swiss territory*, one with blue for *is part of the ocean*, and so on). If, and only if, all the layers are true, the map is true (figure 2.5). A semantics is a way of understanding certain representations by looking at what makes them true or false. The proposed semantics allows you to see in which ways maps can be false.

Crucially for the account detailed in this chapter, the corresponding relationship between parts of the map and parts of the territory is a primitive, unexplained notion. And here is where spatial structure matters most. Once you have assigned a certain area of the map to a certain area of the territory, many other assignments flow naturally and automatically in an unstoppable cascade. The actual assignments depend on some precision targets one can set to the map. These precision targets in the norm correspond to the choice of a geographic projection method. Projection methods preserve on the target some, and abstract from other, spatial properties of the projected. The metro map of London famously preserves topology (order and contact)[9] but does not care much about (geometric) similarity and metric properties (shapes and distances). In our account, we set for a very conservative minimalist target: only the preservation of part/whole and topological properties was assumed, while other projections were considered optional as they would have just been a harmless strengthening of the core projection; the more, the merrier, but nothing less than mereotopological preservation could be accepted.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

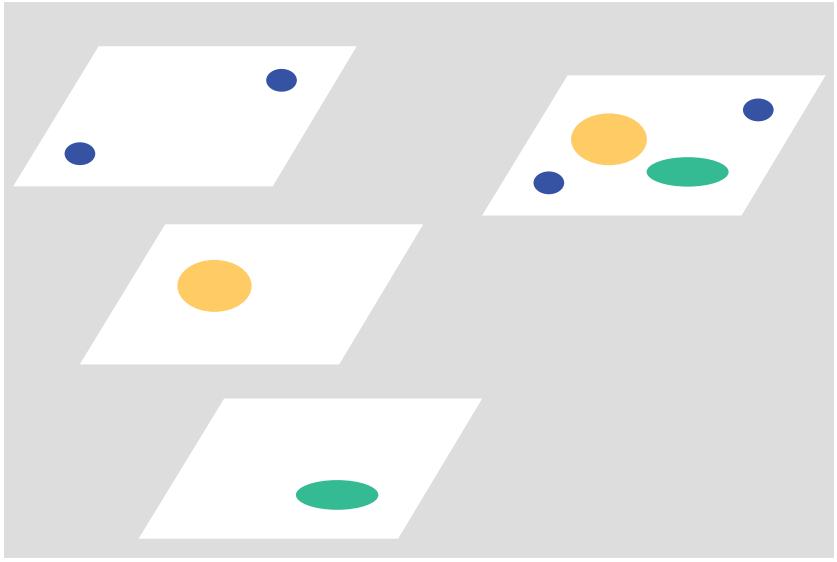

Figure 2.5
Three map layers (left) of a formal molecular map (right) of the same territory. Note that map layers are maximal (if an area is colored blue in the molecular map, it is colored blue in the map layer) and that they support disconnected areas (the sum total of the blue areas in the layer representing woods is what matters). The map layer at the bottom of the pile displays in green the position of woods on the territory; the color white is used here to refer to the space, but it only signals the absence of woods in the territory. Image credit: RC and Achille Varzi.

To see how spatially systematic maps are, let's contrast maps with language. Some singular terms in natural language are referentially structured in a parthood fashion that matches the morphological composition: the singular term "London" is part of the singular term "the county of London," and the referent of "the county of London" includes the referent of "London." However, this pattern does not establish a general case as the referent of "the capital of the UK" does not include, and is rather included in, the referent of "UK" (figure 2.6).

Singular terms on formal maps differ from ordinary language singular terms because they are spatially structured insofar as their subparts and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

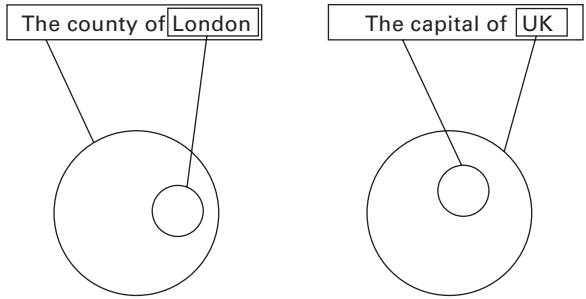

Figure 2.6
The nesting of linguistic singular terms is not necessarily mirrored by the nesting of their referents (the mirroring is the case for maps). Image credit: RC.

super-parts are themselves singular terms that systematically refer to sub-parts or super-parts of the territory referent of the terms. Our account harnesses the iconic structure of regions to justify a whole set of inferences involving parthood and connection of regions in the territory.

Methodology: Formal Maps Are the Normal Form of Maps

I shall explicitly, and sometimes implicitly, consider formal maps as the *normal form* in which we need to put maps to understand their inferential capabilities. Two observations are in order. First, ordinary maps are already quite abstract representations. For this reason, I make the choice of steering the abstraction process toward a point that makes the maps we discuss quite similar to formal maps (as in the instructional map of figure 2.7): systematically organized syntactic areas, with colors expressing properties instantiated in the referents of these areas.

Second, this choice to focus on formal maps is task-dependent, as is any choice of the level of abstraction in using a certain representation. My own task is to elicit two computational strengths of maps, both those related to situated (navigational) inferences and those related to non-situated inferences. These are core computations we perform when using maps that would be impossible, or likely highly impractical, to perform if we did not have maps at our disposal. I could have been interested in other computationally salient features of maps, such as the readability of city names, of coordinate grid structures, and so on, but these choices would

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

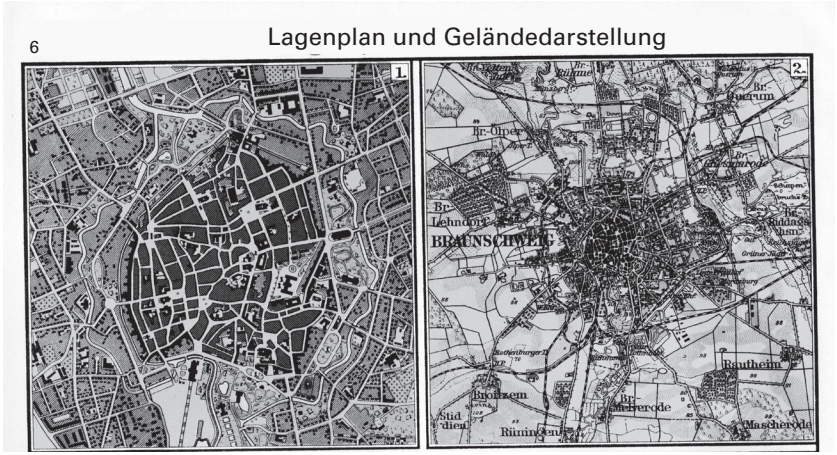

Figure 2.7
Instructional atlases use abstract, "mute" maps to illustrate the process of mapmaking. *Schulatlas*, undated, first half of the twentieth century. Author's collection. Image credit: RC.

have engaged normal forms other than the formal map format that is the focus of this chapter.

Distraction: *How maps lie, how to lie with maps.* A good deal of the literature on maps points out the many interesting aspects of arbitrary choices made by mapmakers and the many ways in which maps have the capacity to lie or present a distorted representation of the territory (figure 2.8) and does so possibly to make other features easier to grasp (Monmonier 1996) or to pursue other agendas. I do not have anything substantive to say about that, except that yes, maps can and do lie and distort and omit. Of course! My point is to state that some structural features of the territory must be preserved by the map not because we want a map to be accurate but rather because we want it to be a map. If Bourgogne is now a proper part of France, nothing we would call a map can represent the territory that is now France as a proper part of the territory that is now Bourgogne. Whatever this representational item is, it is either not a map of that territory or not a map at all.

The Automatic Assignment: Semantically Active Parts

Mapmaking is one thing and the semantics of maps another, although the two of course intersect. Here is an important divide and point of contact. Suppose you want to draw a map of the room you are in. You know that your room is a square, so you start with a square (figure 2.9).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

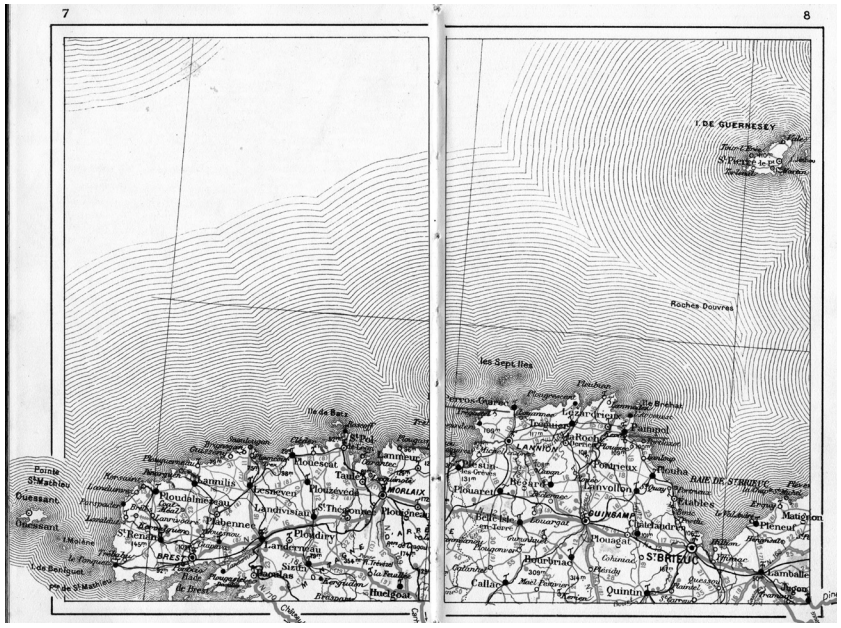

Figure 2.8
For the route map, the sea does not matter, and its representation is filled with purely decorative elements espousing the coastline, which we can call "graphic echoes" of the coastline. The map literally lies about the sea, but interpretive practices prevent us from using it for sailing or underwater prospection. *Carte Ford* 1932. Author's collection.

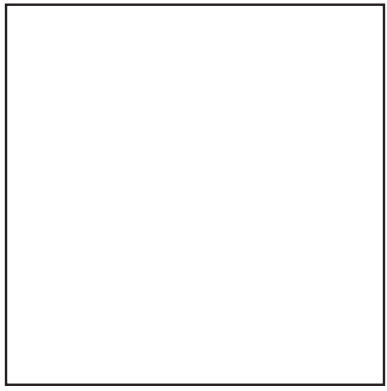

Figure 2.9
Your room.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Mind this step! As soon as you draw a square on the map and orient it, the area inside the square instantly becomes semantically "active." Parts of it *already* and *automatically* refer to parts of the room, even if you are not able yet to precisely tell which one refers to which one.[10] The next step in mapmaking is precisely to find out which parts of the piece of paper refer to which parts of the territory, which is what keeps cartographers busy. But all the parts of the map *already* refer to their referents, no matter what you know about them. If they didn't, you could not make mistakes in mapping. And if they didn't, some claims about mistakes would not even be intelligible today. Suppose travelers were mapping an uncharted part of the Asian continent in the seventeenth century. After they finished, one of them would exclaim: "We were wrong, this area here in the middle actually corresponds to Ireland!"

Distraction: Did I forget to mention the choice of a projection method? Yes, please choose a projection method. Which map regions correspond to territory regions are a consequence of your choice of a projection method and of the automaticity of the assignment. This will be clear in the following example.

Assignment: An Episode of Recycling

Approaching the small French island of St.Pierre-et-Miquelon off the coast of Canada with a broken tiller, French sailor Eric Tabarly is engulfed in fog (Tabarly 1974). He does not have a map of the island and makes do with a portolan giving the latitude and longitude coordinates of salient points (lighthouse, rocks, shoals, and so on). To minimize error, he "recycles" a map of the Belle-Ile area on the Brittany coast, which is drawn for the same latitude. (Consider: the area comprised between two close-enough meridians and two close-enough parallels resembles more and more a rectangle when we approach the equator, more and more a triangle when we approach the poles.) On that map, he places marks for the salient features registered on the pilot book, keeping a constant latitude and working out a transform for the longitude. The rough map (figure 2.10) allows him to both have a visual impression of the approach and make graphical route calculations.

The recycling of the map of Belle-Ile comes about by dereferencing the original assignment and by assuming that the vehicle now represents points in the St.Pierre-et-Miquelon territory. But at the beginning of his remapping, Tabarly does not yet know which points on the map correspond to

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

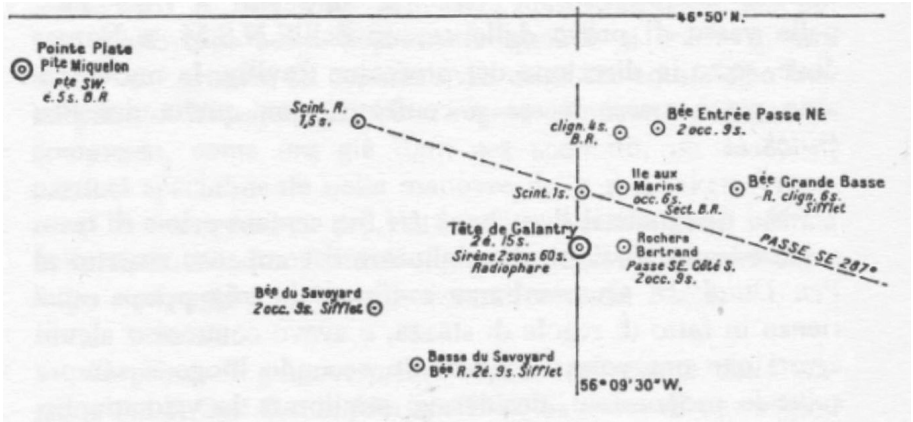

Figure 2.10
Eric Tabarly plans his route by recycling the grid of a map at the same latitude. The text accompanying this image reads: "This chart used to enter the port of Saint-Pierre was made from a map of Belle-Île at the same latitude. I could thus have had the mercatorian grid and the scale without having to draw them. The Tête de Galantry is the southeast tip of the island of Saint-Pierre. The upper lighthouse of the alignment is located ashore in the port of Saint-Pierre. The lower lighthouse is located at sea on a rock."

which points on the territory. Little by little, he discovers corresponding points by referring to the pilot book. I think we can respect this intuition, that in mapping the referents of the map areas are not assigned but rather *discovered*.

Mapmaking and Discovering a Territory's Structure

The holding of a referential relationship between parts of the map and parts of the territory is a metaphysical datum; things are just like that. But in addition, we *know* that this relationship holds. However, our knowing or not knowing *which* parts of the map refer to *which* parts of the territory is a further epistemological problem. Knowing that the parts already refer to some parts of the territory once the map space is cast does not make you know which parts of the map refer to which parts of the territory. And this further piece of knowledge is not necessary for knowing that there is a relation in the metaphysical sense. The metaphysical datum is not sufficient for the epistemological achievement, and the latter is not necessary for the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

former. The converse statement also holds. Knowledge of the fact that a certain map region refers to a determinate place is sufficient to make the claim that there is a referential relationship, and the existence of a referential relationship is necessary to claim that this particular map region refers to a particular region in the territory. (If there was no such referential relation, then there would be no knowledge thereof.) Metaphorically speaking, the work of a cartographer is like scratching the scratch-off layer of a lottery ticket and uncovering correspondences place by place. (Of course, cartography in practice is much more difficult than that.)

The metaphysics of the referential relationship between maps and territory intersects mapmaking at another point. In a map, the mapmaker sets a certain projection method, and what map regions correspond to what regions in the territory depend on this choice of projection. But once the choice is made, no room is left for arbitrariness.[11] And, as argued in the previous sections, there is a kind of bottom line here: some assignments, specifically those that do not preserve the basic mereotopological relations, are map-destructive. They are map-destructive in the sense that the item you end up with is not a map in any recognizable sense.

There Are No Gappy, Idle, Glutty, or Ambiguous Parts of a Map

What other things do we know about the automatic assignment, besides knowledge of what parts of the map exactly refer to what parts of the territory?

- We know that the map *is not gappy*: there is no area of the territory that is not referred to by an area of the map.

Important caveat: This principle and the following ones concern the *assignment* of territory regions to map regions; in a completely different sense, maps can for instance be gappy, that is, omit to represent certain features of the territory. To make the idea reasonably vivid, a maritime map may fail to represent an island. (This is either because it is not yet known or because inhabitants or explorers have reasons to hide it from you, because of this and that.) But it cannot fail to represent the *region* where the island is situated.

- We know that it *is not idle*: there is no area of the map that does not refer to an area of the territory.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

- We know that it *is not glutty*: no two different areas of the map refer to the same portion of the territory, unless they are replicas (e.g., when you find a region both at the leftmost and rightmost extremes of a world map; see figure 2.1 where Europe is represented twice in two different portions of the same map).
- We know that it *is not ambiguous*: one and the same part of the map does not refer to two different territories (figure 2.11).

These are the key elements of our knowledge of what a map is, and together they articulate a core *intuitive theory of region assignment to maps*.

The automatic assignment is a crude fact about maps, and the four principles flow from our knowledge of it. Knowing which parts of the map correspond to which parts is an epistemological endeavor. But *knowing that the assignment is automatic* is part and parcel of our understanding of maps. Items that did not satisfy the four requirements would be very hard to use for navigation. This brings us to an important set of claims. But first let us describe natural generativity.

Natural Generativity

Some representational artifacts are *naturally generative*. The notion has been introduced by Flint Schier (1986; see Casati and Tappolet [1987]) and can be summarized as follows. Some representational systems (e.g., pictorial systems of which photographs are an example) are such that learning to interpret a few examples of them is enough for you to be able to interpret all examples, provided you are *perceptually familiar* with what they represent. Take, for example, spoken language as a contrast class. If you learn and understand what "My cat attacked your dog" means and are perceptually familiar with scenes of domestic animals fighting, you are not *thereby* in the position to understand what "My doll smiled at your car" means, even though you may be perceptually familiar with dolls, smiles, and cars. What you definitely need is a knowledge of the meaning of "doll," "smiling," and "car." Perceptual familiarity with the referents is neither necessary nor sufficient for interpreting most natural language utterances (Distraction: except for marginal cases of onomatopoeia). On the other hand, perceptual acquaintance with physical items is both necessary and sufficient to interpret *pictures* of those items, once you get a representational epiphany by

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

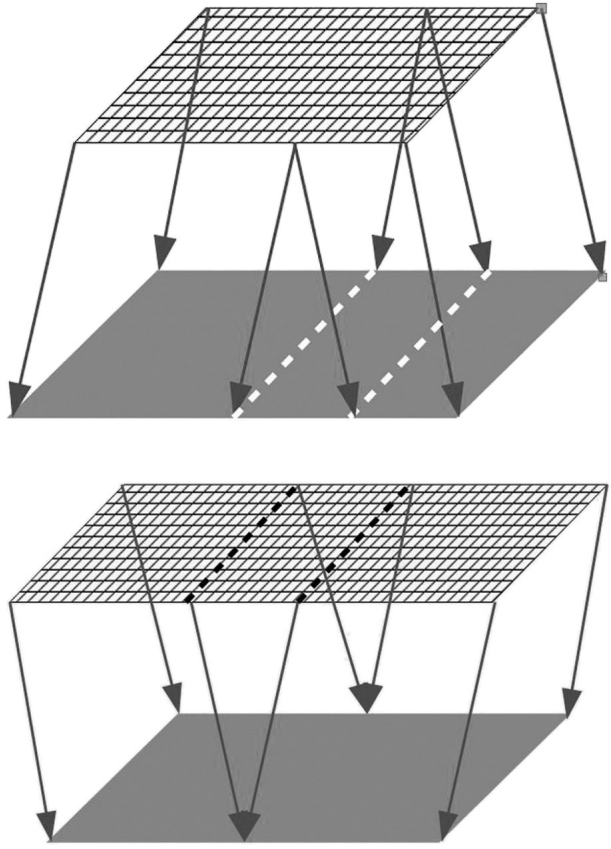

Figure 2.11
What we know intuitively about the assignment structure of maps: that it cannot be (a) gappy, (b) idle, (c) glutty, or (d) ambiguous. Image credit: RC.

looking at a picture. If you understand what the picture of a cat attacking a dog represents, perceptual acquaintance with dolls, smiles, and cars is all you need to understand the picture of a doll smiling at a car.

As we shall see, natural generativity can take on different forms that depend on the type of acquaintance required. Perceptual acquaintance is one such type. In the case of maps, less than perceptual acquaintance seems to be required. Rather, maps require that one possess a minimal spatial competence.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

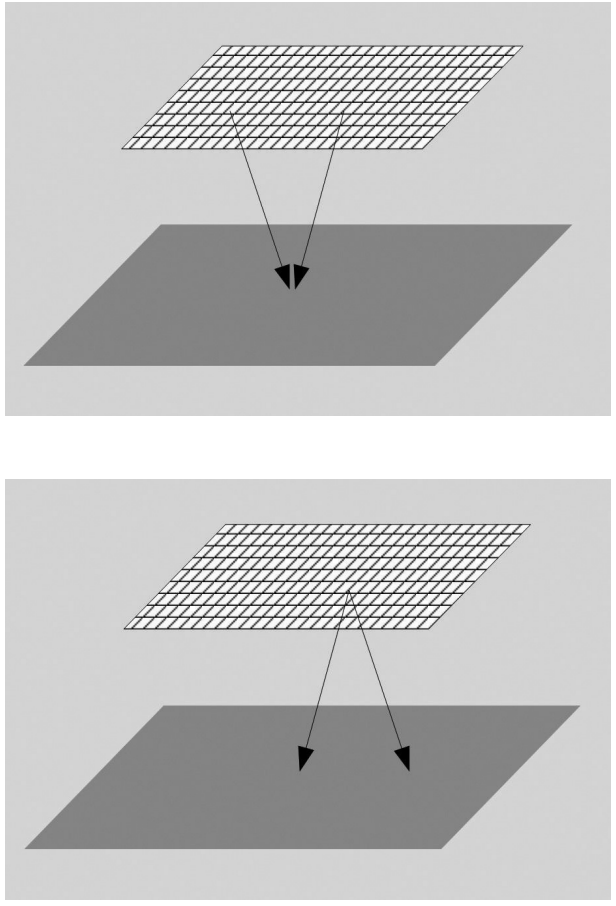

Figure 2.11
(continued)

Maps Are Hypergenerative

Minimal spatial competence about maps dovetails with the hypergenerativity of maps. Consider the following scenario. A child is playing with a short red woodblock and a tall green woodblock. She lays both on the table and asserts: "This is Fido, and this is Tweety" (these names being the actual names of her pets). As soon as the two blocks are laid on the table and names are distributed, an instantaneous understanding of the spatial relationship of the two blocks as representing a spatial relationship between Fido and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Tweety follows. Say the child moves the two blocks close to one another and puts them in contact. We immediately understand this movement as representative of a change in the spatial (both metric and topological) situation between the referents of the two blocks. The setting (blocks + table) is automatically read as a map. I call this type of phenomenon the *hypergenerativity* of maps.

Pictures are naturally generative because to understand a few examples of pictures is enough to understand most new instances, provided we have perceptual acquaintance with the referents of the latter. Maps *hyper*generate because not even perceptual acquaintance with types or tokens of their referents is necessary. At this point, we can see that there is an important nexus between hypergenerativity and the automaticity of the assignment. In the end, hypergenerativity *is implicit knowledge of the existence* of an automatic assignment.

When we hypergenerate, we rely on the activation of mental spatial representations rather than the activation of perceptual recognition. Contrast again hypergenerativity with the contribution of perceptual recognition to natural generativity in pictures. In the case of maps, you do not need to be acquainted with the exact *spatial situation* that involves the objects you are mapping in order to understand a map. Though maps can be very rough, they can still be correctly interpreted, even by young children (Winkler-Rhoades et al. 2013). The natural generativity of pictures and the hypergenerativity of maps facilitate learning of pictures and maps, respectively, without disposing of it or of its necessity. With time, we get more and more proficient at using and interpreting both maps (including dynamic GPS displays) and various types of pictures (e.g., drawings, cartoons, and movies), and we consolidate, M1-style, some aspects of these abilities.

Hypergenerativity goes hand in hand with roughness. The appearance of the two blocks can be far removed from the appearance of the two pets. What matters is that the representational vehicles we use take up space and are situated in places. Their being red or tall does not matter. (It can be made to matter, but it is not necessary. In our example, the tall green woodblock refers to a smaller fluffy yellow bird, and the smaller red woodblock refers to a larger furry brown dog.) The iconic pictorial properties of the representational vehicle are orthogonal to the spatial properties of the map. Blue can represent roofs on a map of a town, even if all the roofs in the town are red. Roughness makes simple maps very similar to formal maps in the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

above-defined sense. Formal maps, as described previously, are the more abstract possible maps that do not contain representational gaps.

The Generativity of Maps Is Specific

The hypergenerativity of maps is immediately available in the two colored blocks + table situation. This is an instance of a *representational epiphany* that is *map-specific*. You immediately grasp that a certain spatial situation (the spatial relation between Fido and Tweety) is signaled when the two blocks are placed on the table. Spatial competence and imaginative rather than perceptual capabilities are involved here. When spatial enabling is provided by the props, it is not that the table on which the two blocks are laid is "recognized" as being a representation of a plaza or as a meadow. Instead, the table is just taken to be an enabling spatial field for imaginative play. The two wooden blocks denoting Fido and Tweety do not trigger any recognitional ability in you, and if they do, either that does not seem to matter or we can make the blocks as abstract as needed so as to eliminate any recognitional trigger. The only thing to be preserved is the fact that the blocks occupy space and can be located in that space. The only space that is relevant for map hypergenerativity is the surface on which they are placed.

That perceptual knowledge of the referent is not the condition at stake is implied by the fact that perception is perspectival, and therefore supports occlusion and is potentially silent about occluded areas, whereas map competence is non-perspectival and, hence, does not support occlusion and cannot be silent about any of the represented areas or their footprints.[12]

How deep can roughness go? How much can we abstract, omit from maps? It depends. What are the inferences we want to account for? As we have seen, elevation maps only measure the elevation of the footprint of buildings because we are interested in, say, watersheds and not the physical buildings themselves.

Map's Generosity: Eased Abstraction Is Supported by the Automaticity of Structured Assignments

Many threads converge at this point. A great cognitive advantage of maps is that they support very partial, abstract, or selective information, which in turn makes it easier to compare how certain properties are distributed over

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

a territory. As Tyner (2014) observes by contrasting satellite views and maps of the same territory, maps can omit many items satellite photographs are forced to display. Two styles of abstraction are actually possible: the display of only some properties (e.g., "blue = water-covered areas") and the display of only some distributions of a property (e.g., "turquoise = lakes with an area of more than one square kilometer"). Maps can endure a lot of stripping out of displayed properties; no matter how much one selects, in the end, the set of map regions—even empty regions—that make automatically reference to some parts of the territory, will survive.

This has consequences for how we use maps. One suggestion is that we make-believe navigate *in* them. Make-believe (Walton 1990) comes in kinds. Our engagement with pictures is recognitional make-believe. You imagine being in front of and perceive the scene the picture represents because your recognitional abilities are triggered by the picture. Once more, we see that recognition is not essential to understanding maps. In fact, you may never see a marsh area from an aerial view, but you understand that what you have at hand is a map of the marsh area you are trying to cross. Conversely, certain types of make-believe are specific to maps. One type used for maps is navigational make-believe. You put your finger or placeholder on the map, and you walk it around. Map-specific make-believe practices include the ability to imagine being a giant, or even a bird-like flying entity, who is looking at the actual territory when looking at the map.

Maps Are Inference-Friendly—As a Byproduct of Their Format

What do we use maps for? We use maps to find out facts about the territories they depict, in particular places in the territory that are not visually accessible from where we are, or that are not visually accessible at all. By looking at a map, we can find out that Belgium is smaller than France, that Italy does not share a border with India, that the United States is not a contiguous territory, and so on. Some of these facts are plainly coded in the map. Other facts are available through a bit of inferential work, and this depends on the particular query one asks to the map. Moreover, some of the relevant inferences drawn from maps are non-situated, whereas others are situated. Take, for example, the following non-situated inference. To go from the capital of Switzerland to the seashore, one must first walk, bike, take a train, or fly through one from a series of countries (Italy, France, Germany, Liechtenstein, or Austria). Conversely, compare the following

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

situated inference. To go from where you are, at the foot of the Eiffel Tower in Paris, to a specified nonvisible target, the Pont d'Austerlitz, you must walk a bit north, reach the Seine riverbanks, follow the Seine upstream, and stop at the fourteenth bridge.

The important cognitive advantage of maps over verbal descriptions of a territory is that this kind of fact-extraction work is relatively cheap. In most cases, visual inspection is sufficient for fact extraction once an appropriate query has been framed. Maps are navigation-friendly. A territory's representation is spatially organized in a way *that allows us to navigate it as if we navigated the territory itself*. Maps are also comparison-friendly. Many of a territory's spatial features are preserved on an associated map precisely because the latter is a map; and humans are good at the visual inspection, classification, and comparison of spatial features (as seen in the example of the timetable in chapter 1).

But these cognitive advantages are just by-products of the spatial features of maps; *they come for free* given that maps are by their essence spatial structures and that they have the same dimensionality as the territory they map. Other visualizations, such as diagrams, *inject* spatial features into the representational vehicle to make extraction cheaper. There is no particular reason, for example, to use *spatial* variation to represent with a diagram variation in the average body mass of the Norwegian population. Most of the current discussion can be read as a tribute to the physicality of maps, the fact that maps are spatial objects unfolding in space. The abstract features of maps are constrained by their physicality, and this in turn enhances levels of extractability and potential inferences.

Of course, there exist hybrid representations that combine the syntactic and semantic features of maps with the syntactic and semantic features of diagrams. Functional magnetic resonance imaging (fMRI) visualizations (statistical maps) are one example of this combination. Maps of the world with superimposed data columns for population are another. However, these remain maps with added diagrams.

Camp (2018) objects to our account that the use of conjunction as the only posited compositional operator creates a computational challenge to map users: if blue stands for water within a certain region on layer 1, if green stands for woods within a certain region on layer 2, then to ascertain whether woods border water would require a lot of inferential work in moving from one map layer to the next. We would need to determine where blue is, determine where green is, and then relate the regions

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

under discussion identified. In our semantics understanding of maps, it is contingent on layers being printed on top of each other that we can easily retrieve the relationships between water and woods. In and of itself, the only thing that our semantics can see are layers and conjunctions of layers.

Camp's objection is well taken. However, we locate the difference at the computational level, not at the informational level. A set of layers and a superposition of the items in a set of layers are informationally equivalent. The only difference between the two is computational. Considering maps as layers of transparent sheets (each of them a map layer, printed with one single color) may help reconcile the intuition that maps are computationally friendly even if conjunction (as a combination of many single-colored map layers) is the only compositional structure they have.

Back to Ordinary Maps

Our discussion of maps has been methodologically reduced to a discussion of the theoretical notion of a formal map. This discussion has presented formal maps as the normal form of any map, relative to the task of explaining core computational inferences based on maps. It is thus important to show to what extent ordinary maps can be put in normal form, that is, translated into formal maps.

Steps to the Translation

Figure 2.12 displays an ordinary, if simplified, map of a territory. Many types of symbols appear on the map: regions, colors, boundaries, and singular terms (names of places). Translating it into a formal map requires making some choices. Figures 2.13 and 2.14 show two possible ways of proceeding.

Types of Symbols on Ordinary Maps

What is the status of other *types* of symbols on ordinary maps besides the now-familiar symbols of region and color patch type? For instance, are icons (e.g., for trees) and names of localities ("Paris") of the same type as colored regions? Are they merely of different orthographic type, the way "Paris" and "PARIS" are, or do they introduce genuinely different semantic ingredients?

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

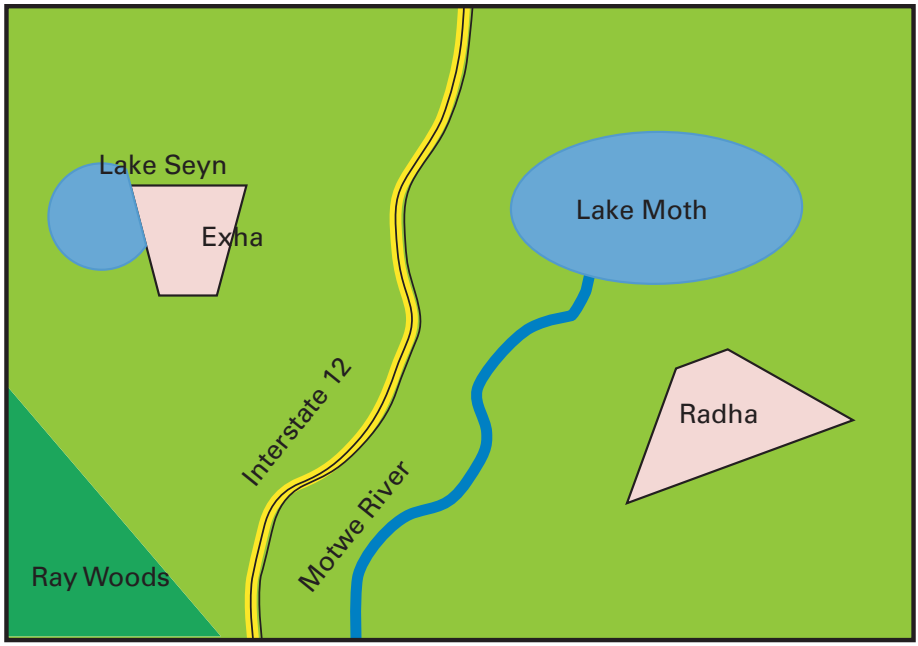

Figure 2.12
An ordinary simplified map of a territory displaying two cities (Exha and Radha), cultivated land, woods, two lakes, a river, and a highway. The map is already quite simplified insofar as it does not display icons, isolines, and so on. Image credit: RC.

The answer, in some revisionary fashion, is that icons and names are all the same syntactic type: the type of *color patches*, that is, predicative expressions. They differ only orthographically from one another. They are of different orthographic types. That different orthographic types correspond to one and the same syntactic type and to one and the same semantic type is just a mark of the difference between ordinary and formal maps. In the same sense, different morphological types in ordinary language may on some accounts be assigned the same semantic representation ("John *is tired*" or "John *stinks*" if, for example, we bar event-based semantics).

There Are No Other Singular Terms on Maps Except Map Regions

To see the issue of symbols on maps in greater detail, consider the fact that there are, of course, names (i.e., of cities, countries, and so on) on ordinary

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 2.13
Two translations of the map of figure 2.12 into formal maps whose actual content is determined by the cartographer's choices. Note that boundaries and written names have disappeared and only colored regions are on the map. The map on the right considers "being the city of Exha" and "being the city of Radha" different properties (same for "being Lake Seyn" and "being Lake Moth"). Image credit: RC.

maps, such as "Paris" or "France," that one would naturally consider *not* as predicates but rather as singular terms. The idea defended here is that these linguistic items are not used as referential devices on a map in the way in which they might be used in a sentence. Hence, their proper translation on a formal map should be a color, not a region. In a sentence such as "Paris is beautiful," "Paris" refers to the subject of the predicate ". . . is beautiful," that is, the town Paris. But on a map, the sign "Paris" does not directly add to the map's referential powers. *Reference* to the region where Paris is located *is already made by the relevant map region*. I offer two readings for the occurrence of "Paris" on an ordinary map. In both readings, "Paris" is a component of a predicate. It might be a component of the predicate ". . . is named 'Paris'" or of the predicate ". . . is identical with (a part of) Paris." Insofar as it is a component of such predicates, the value of "Paris" is either

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 2.14
Two map layers (or stages) of the right-hand map of figure 2.13. Each of them displays the spatial distribution of one single property and also represents the area where the given property is not instantiated. For instance, whatever is not green in the left map layer represents an area that is not cultivated land. Maps such as the right one in figure 2.13 are superpositions of map layers, acting like layered transparent sheets. Each layer is true when both the area in the territory that corresponds to the color for a certain property has the property in question and the area of the territory that corresponds to the blank area does not have the property. The map is true whenever all of its individual layers are true. Image credit: RC.

Paris's name, or Paris. The predicates ". . . is named 'Paris'" or ". . . is identical with (a part of) Paris" will then be given an unsurprising translation into a color. This would be orthographically clumsy on an ordinary map.

Entire chapters of textbooks on cartography (e.g., Tyner [2014]) are devoted to syntactic elements other than colors, such as symbols, icons, and text. As I repeatedly said, it would be impractical to have maps without these syntactic elements, and formal maps are not proposed here as substitutes for the types of maps that we find in bookshops. Scholars such as Pratt (1993) and Camp (2018) have followed the path of providing a semantics for syntactic elements other than colors. For instance, Camp's proposal has two parts, one for what she calls marker types and one for marked-zone types. Marker types are icons, words, or pictograms; they are placed on

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

relevant parts of a map (e.g., a pictogram for lavatories 🚻) but "their non-locational spatial properties are not projected" (the lavatory pictogram has the shape of a man on the left and of a woman on the right, but the actual lavatory does not need to have a male section to the left and a female section to the right). In our account, marker types are not present. They are translated, rather, into colors and colored regions, and so we do not need to deal with them in the discussion. All we need to make sure is that we are able to translate marker types into colors.

For Camp, marked-zone types are lines, shapes, and so on, corresponding to objects in a territory. Their spatial properties are projected. The shape of the zone corresponding to Lake Moth (cf. figure 2.12) projects as the shape of the lake. Marked-zone types are images (silhouettes) of the objects they represent. Maps are true if the items denoted by marked-zone types cover areas within a given region in such a way that topology is preserved. The main difference from our approach is ultimately a difference in ontology. Camp endorses as if it were a "Leibnitian" ontology according to which space is the by-product of objects and relations between objects; we endorse a "Newtonian" ontology that also allows for regions and relations between regions. (How do we choose? A neutrality hypothesis here is that the two ontologies are at a minimum operationally equivalent, and that the choice of a region-based or an object-based system is more a matter of ontological taste. A less neutral hypothesis: the Leibnitian account makes questions about empty zones on maps a bit more urgent. For us, empty zones are semantically active, even before any discussion of the absence intuition. Empty zones refer to regions in a territory, even if they are not marked or colored [chapter 2]. They are part of the map's syntax. These zones create perhaps a challenge to Camp's account, both informationally and computationally.)

Things That Are Formal Maps That We May Not Want to Consider Ordinary Maps

The principles used to characterize formal maps turn some intuitively awkward things into (close approximations of) formal maps. Every map that results from a mereotopology-preserving projection of a formal map is a formal map. For instance, a mirror-reversed projection of a formal map of North America (figure 2.15) that preserves north and south but swaps east

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

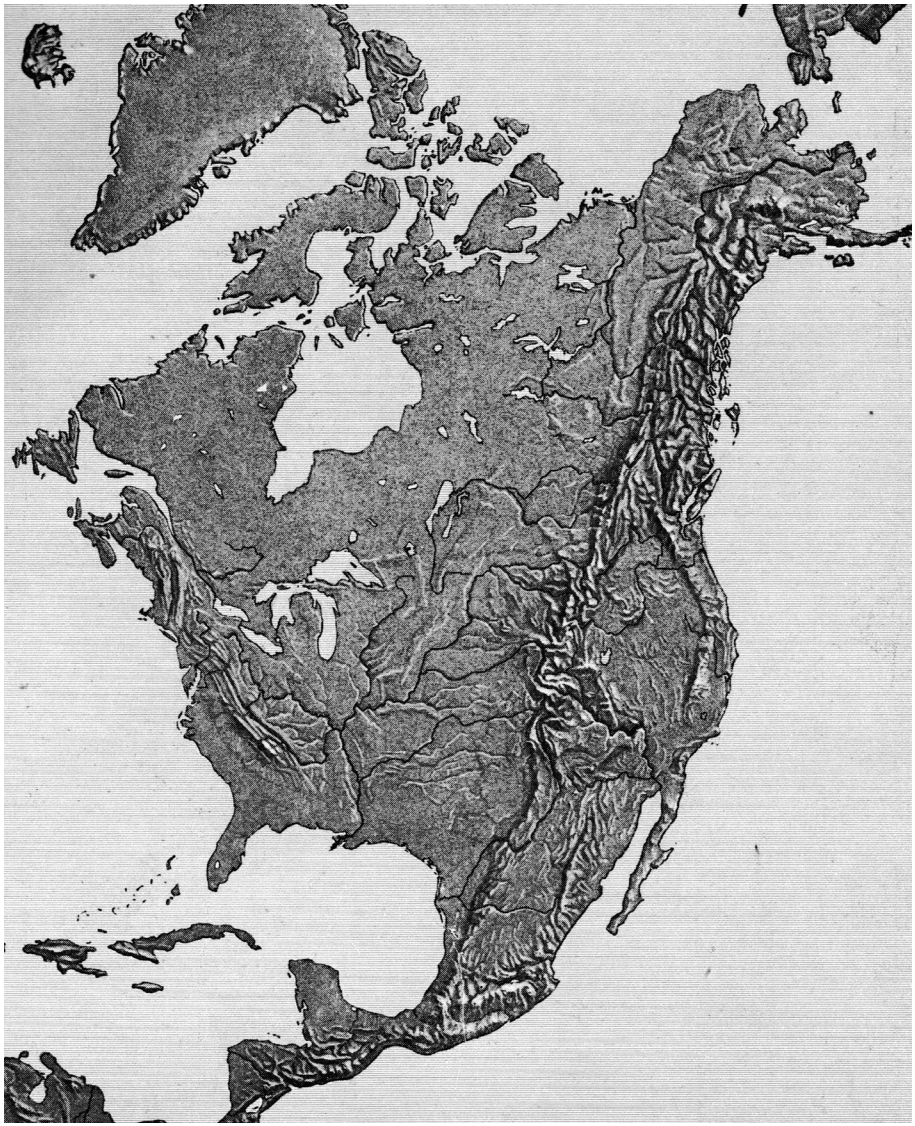

Figure 2.15
An east-west inverted map of North America. Source: *Atlas classique*. Paris: Hachette, 1931; original image has been flipped horizontally by the author. Author's collection. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

and west is a formal map of North America. And so is a stretched formal map of North America.

> Distractions to the effect that an east-west inverted map (figure 2.15) of North America is not a map of North America can be easily addressed (Casati 2013). An east-west reversed map of North America presents North America as seen from the center of the Earth. Such maps have a use. For instance, they can be used for matching the longitudes and latitudes of cities on Earth with the celestial longitudes and latitudes of stars.[13]

Overlapping Colors

Colors can overlap on a formal map, whereas only some ordinary maps allow for color overlap (e.g., maps of the geographic distribution of bio taxa). Coloring on a formal map is not subjected to any particular constraint. The only nonoverlapping constraint, trivially, concerns a color and its absence. Different colors can compete for the same region. Color competition between regions only poses a technical, graphical problem, as when we use superposed transparencies or grids. The nature of the represented properties can, of course, dictate further constraints, but these do not belong to the overall formal map structure. For instance, on an elevation map, the color corresponding to "altitude between 1,000 m and 2,000 m" and the color corresponding to "altitude between 2,000 m and 3,000 m" will never overlap, as a matter of necessary fact. The fact that isolines cannot cross on a map is not a syntactic rule.[14]

Derived Symbols and Redundant Symbols, and the Four-Color Theorem

Ordinary maps employ redundant symbols. For instance, provided we use enough colors, boundary symbols become redundant on maps, in the same way that vertical lines used to delimit bars are redundant in music scores (see chapter 7). Boundary lines both computationally and visually supervene on the color difference between adjacent regions. This being said, we may still want to draw boundaries nevertheless. Thus, what is their status on a map? These boundaries are not proper symbols in the same way that the color difference between the two regions is not itself a symbol (it is not part of the map lexicon). Formal maps as described here, in fact, do not explicitly represent boundaries. Taking a different approach, one might want to skip colors altogether and use boundaries to signal the relevant demarcations.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Distraction: The four-color theorem (stating that "any map in a plane can be colored using four colors in such a way that regions sharing a common boundary [other than a single point] do not share the same color") is not relevant for formal maps. This is because on a formal map—and on an ordinary map, for that matter—color patches can be scattered (non-self-connected). Though the contiguous United States is disconnected from Alaska and Hawaii, a single color is used for all of the US states. In our context, the four-color theorem (and similar constraints) is only moderately relevant. It is a condition on ordinary map design that has nothing to do with the use of colors as predicates. A map where Japan and France have the same color may well conform to the four-color theorem, but it will also express the fact that Japan and France share a property.

Dissective Properties

If the color blue expresses the property of *being water*, then any territory region that is referred to by a map region that is covered with a blue patch on a map will be a watery region. Of course, this works fine with water, for water has watery parts. Being water is a *dissective* property—every part of something that has the property of being water (down to a certain size) has the property of being water (Nicolas 2018). And a blue color patch is particularly well-suited for representing a dissective property such as being water because being a blue-colored patch is a dissective property itself: every part of a blue-colored patch is a blue-colored patch. But not all properties are dissective, and using colors for representing non-dissective properties has its shortcomings. If blue expresses the property of *being a lake area*, then non-lake parts of the territory end up being described as containing a lake because of the dissectivity of blue. A simple fix can consist in limiting the type of expressible properties. Red, for example, will not represent the property of being identical with France, but the property of being identical with a (possibly improper) part of France. But surely formal maps encounter here an expressive limitation. One may conclude that formal maps can only see dissective properties in the territory. They force a kind of "grinding" of territorial properties into uniformly spread properties (Pelletier 1975).

Geographic Maps and Time

Related to water, here is a side but somewhat interesting issue. Take a cell (a space creating a square between grid lines) where only the sea is found

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

within the represented territory of the 1809 Abū Qīr map by Pierre Jacotin (figure 2.16).

What exactly does the Abū Qīr map represent in this chosen square? Sea? Some sea? An expanse of water? (Consider that probably no single molecule of water [or arguably very few] that was there when the map was designed was still there at the time of printing.) More generally, what do the sea portions of maritime maps represent? We do not want maps to be only instant-true, although most maps are time-bound (maritime maps come with a lot of timestamps and warnings about the necessity to use the latest, most updated version).

Two Routes to Information, Two Inferential Paths

Another limitation of formal maps concerns the expressibility of spatial relations. Suppose (figure 2.17) you have three colinear cities *a*, *b*, and *c*, with *b* between *a* and *c*.

As it is depicted in figure 2.17, the distance *ab* is smaller than the distance *ac*. In a sense, one can observe that a system of distances on the map has the power to represent a system of distances in the world (it might be held that the former covaries with the latter). Still, distances are not colors on a map, and thus are not map symbols. Here, formal maps appear to be limited in their expressive power. As a matter of fact, they seem to be able to express monadic properties only (i.e., *being watery*). If they purport to express a relational property (*being larger than x*), they do so by absorbing the relational component into the predicate. *Being between 1,000 and 2,000 meters* is a monadic property into which a number of relational facts are packed.

In a similar vein, by inspecting a metric-preserving formal map of the United States, you can read off that a certain world region is Wyoming and that such a world region is approximately rectangular. The semantics provided for formal maps accounts only for the first type of information: a certain color expresses the property of being Wyoming, and that color is distributed over a map region. The world region referred to by the map region is supposed to have the property corresponding to the color. On the other hand, how can we recover information that the world region in question is rectangular? Being rectangular is a property of that region, but surely we would not introduce another color that corresponds to the property of

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

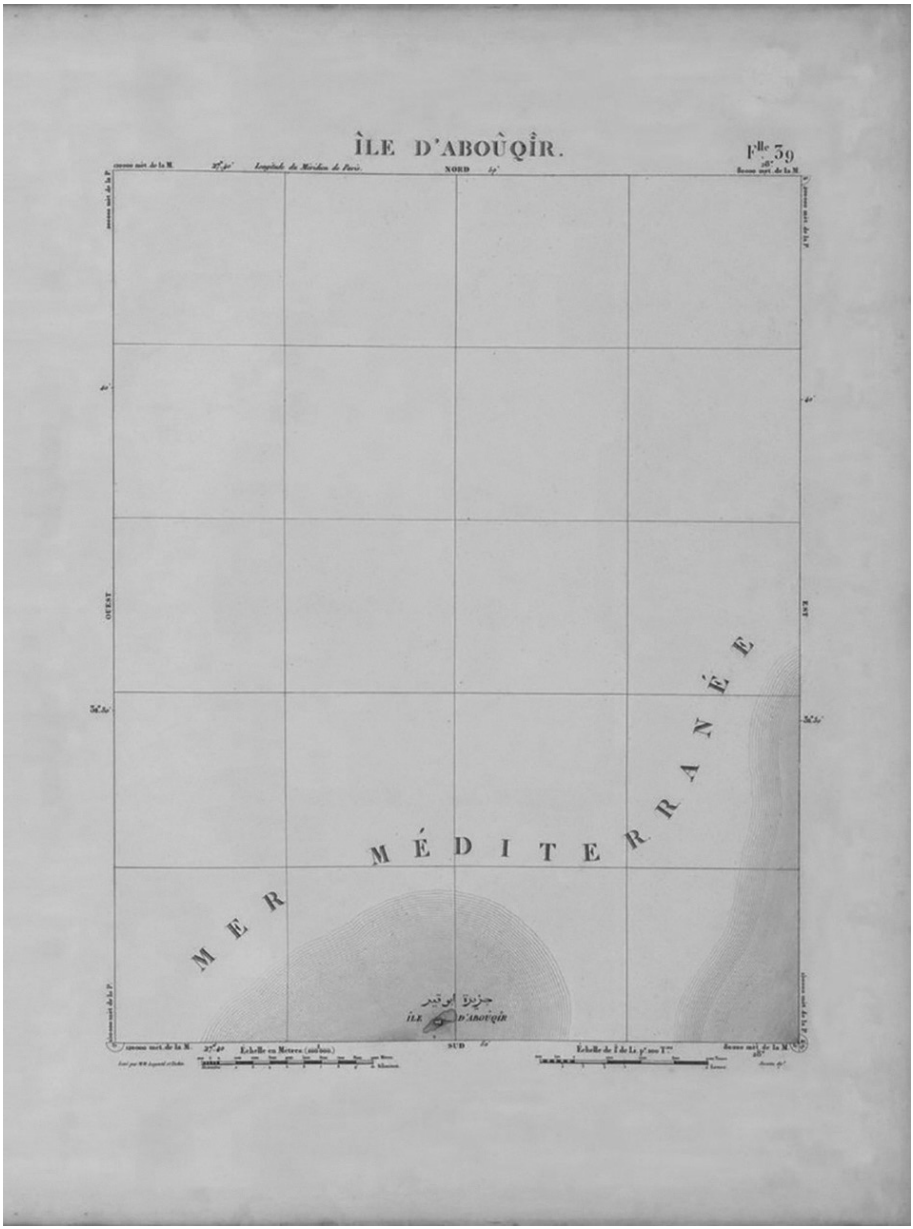

Figure 2.16
Île d'Abū Qīr. Plate 39 of Pierre Jacotin, Description de l'Egypte, ou, Recueil des observations et des recherches qui ont été faites en Egypte pendant l'expédition de l'armée française. *Atlas géographique*. Author's collection. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 2.17
Three cities. Image credit: RC.

being rectangular. Such a color would be informationally redundant (this is a welcome consequence of the fact that in maps we represent spatial territories by spatial vehicles).

But there is more to that. Colors, as they are used on maps, work like monadic predicates. But, as we have just seen, a lot of relational information can be extracted from a map. For instance, on a map of the United States, we see that Wyoming is not contiguous with Maine. Even so, it is possible to pack relational facts into a monadic predicate. In theory, we could assign to a new color the property of not being contiguous with Maine. However, in this case, some of the obvious inferences from statements that use relational predicates could not be extracted from monadic predicates. If Wyoming is not contiguous with Maine, then Maine is not contiguous with Wyoming. And yet, how can this be seen by comparing the color of Wyoming with the additional color introduced for expressing the property of not being contiguous with Wyoming? Wyoming and Maine would just be painted with two different colors, with no logical relation between them.

The semantic engine for the extraction of this type of spatial information must reside elsewhere. Luckily, we have the structure of the assignment, which is spatial in nature. Thus, we can see how inferential procedures that use maps as premises take two routes. The first route is based on the structure of the assignment. This route is governed by the standard logical rules along with, it can be argued, the axioms for mereology and topology. These license inferences that in their non-logical side encode a *theory* of the domain (spatial regions). The second route is based on the propositional information contained within the coloring. This second route is governed by the standard logical rules. The two inferential routes, obviously, can and do intersect at various points (e.g., in the inference from "There is Wyoming at r_1" to "r_2 is American land," where r_2 is part of r_1).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Using Maps for Making Spatial Inferences

Maps are used to draw complex inferences, backed by specific knowledge. A map of isobars (linking points of equal atmospheric pressure) can be used to infer wind direction and speed. Winds on Earth blow from high pressures to low pressures, like a ball rolling down a slope. This step mobilizes both dynamic and spatial intuitions. Due to the Coriolis effect, winds turn clockwise in the northern hemisphere, about 70° to the normal connecting isobars when on the open sea (figure 2.18). When wind blows on your face in the northern hemisphere, the low pressure will be to your right. Wind speed is inversely proportional to the distance between isobars. Sometimes, an abacus is added to isobar maps to help readily calculate wind speed. Insufficient propositional knowledge of these background facts generates unsound inferences, such as misleading explanations of hurricane spins (turning counterclockwise in the northern hemisphere).

Exploring the Map Jungle

In the previous sections, some interconnected tools have been presented: systematicity, hypergenerativity, same-dimensionality, a skeletal formal semantics, automatic assignment, and the small intuitive theory of maps. Based on these tools, we can broaden our theoretical horizon and begin to look at what items qualify as good candidates for the title of "map." In particular, we consider items that may not usually be called "maps," as well as items called "maps" that are not maps at all. Next, I provide a list to exemplify the kinds of items I have in mind. Some of these items are discussed at a later point in the next sections and chapters (and not necessarily in the order presented here).

First, *items that are maps (even if in some cases we do not usually label them as maps):* ordinary maps (e.g., a map of the Upper West Side you can buy at a kiosk); Borgesian maps (e.g., maps that are superposed onto the territory they map); ephemeral onscreen maps, for example, displays of navigator assistants (a piece of GIS software generates the visual display, and the map is what you see on the screen); models—at least certain types of physical models (chapter 5); fMRI "images" mapping a state of the brain, nano-"images" representing the microscopic structure of materials, radar and sonar "images"; sound reproductions on records (chapter 3); ordered

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

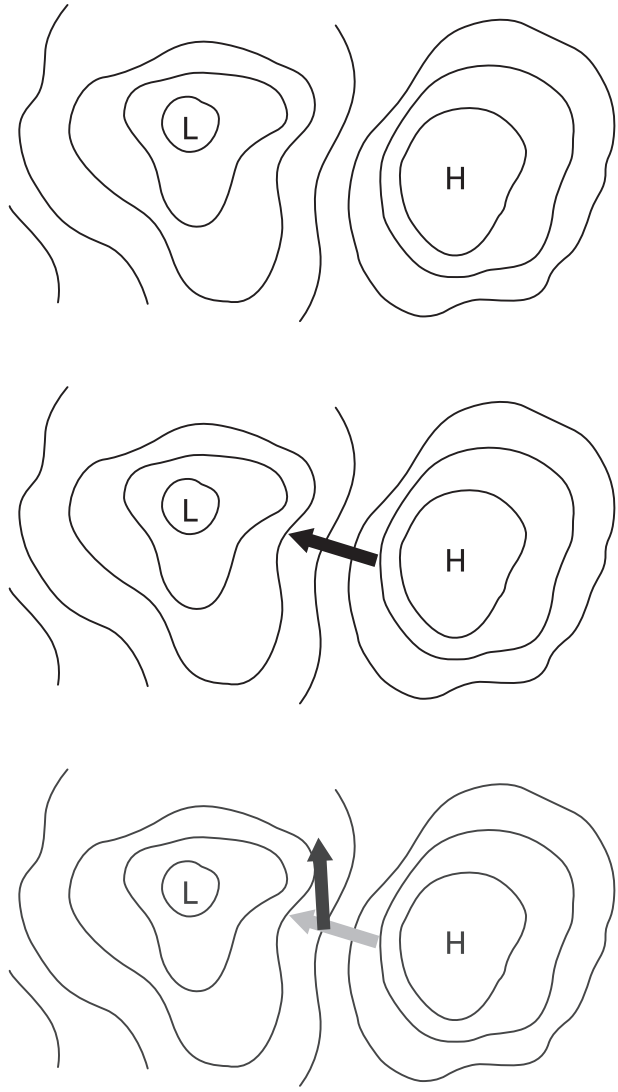

Figure 2.18
Background knowledge used to infer local direction of the wind from a map of atmospheric pressure. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

lists (e.g., the table of contents of this book, or [some] intercom button displays); clockfaces (chapter 6); music notation; certain music visualizations; some aspects of writing (chapter 7); queues (e.g., physical lines of people accessing a service [chapter 8]); and natural maps, for example, star patterns in certain navigational uses (chapter 3).

Second, a few *contentious but interesting cases*: photographs and pictures, including drawings and movie sequences, that represent real scenes (they are maps in a very peculiar sense [chapter 3]); axonometries, somewhere between maps and pictures (chapter 3); quotes (e.g., the singular term "Matterhorn" that refers to a written word, not to a mountain, and whose appearance matches in a map-like fashion the shape of the word); and numeral sequences (they are not exactly maps because they do not encode locations but are sufficiently map-like due to some of their syntactic features) (chapter 8).

Finally, some *items that may look like maps but are not maps*: thermometer graduations (chapter 8); conceptual "maps"; diagrams (most of them); and metro "maps."

In the end, there will be a discussion of maps in the brain, such as retinal "images," retinotopic maps, cortical maps, and mental[15] (spatial) maps, that is, the *internal* representations that research calls "maps."

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

3 Pictures and Maps, Pictures as Maps, and Some Other Maps

This chapter explores the relationships between maps and pictures as representational systems at large and draws some conclusions about other types of representations.

To gain an understanding of the organization of the representational systems involved, let's focus on the maps/pictures contrast. Pictures (such as photographs)[1] and maps look enough alike to make an overview of their commonalities unproblematic. Consider the alternate map/satellite views of the same territory available in cartographic software. It does not take much arguing to see in a map a kind of simplified satellite picture, or in the satellite picture a kind of visually enriched map. Moreover, pictures resemble maps, and differ from diagrams and other types of representations, because there is a certain kind of systematicity in the way they represent spatial items by spatially arranged parts. A spatially organized diagram comparing the average body mass of the respective populations of Norway and Sweden is in this sense very different from both pictures and maps. We can spend a lot of time on the topic, but I propose to take a shortcut and to lay bare right away what I think to be the truth of the matter.[2]

Pictures Are Maps of the (Generalized) Albertian Window

The main claim is: pictures, such as photographs and realistic paintings, *are* indeed maps, but *not* of the things they depict. Rather, pictures are maps of *the light distribution from an optical array on an intercept plane*. This form of map is typically of the Albertian Window. Optical array and the Albertian Window are notions derived from ecological psychology, that is, from an informational account of the perceptual scene based on what a situated eye can retrieve of it (Gibson 1978).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Centered on an eye and reaching it is a pencil of light rays that objects emit, reflect, or diffract toward the eye—the optical array. Each object defines a solid angle whose vertex is in the eye (figure 3.1). A plane can cut the optical array at any distance from the eye—an Albertian Window is any such flat plane that is normal to the line of sight. When they produce a picture, photographers, and even painters, are recording the relevant values of light as they are distributed on the plane. In this sense, pictures *map* the Albertian Window; they are 2D representations of a 2D territory, and they preserve topological and parthood relations of the territory.

Distraction: *Flat vs two-dimensional.* "Albertian Window" is here used as a reference-fixing expression. Windows are, as a norm, flat as well as two-dimensional. An important idea to note is not that the picture maps a completely flat Albertian Window, but that it maps a two-dimensional pattern.

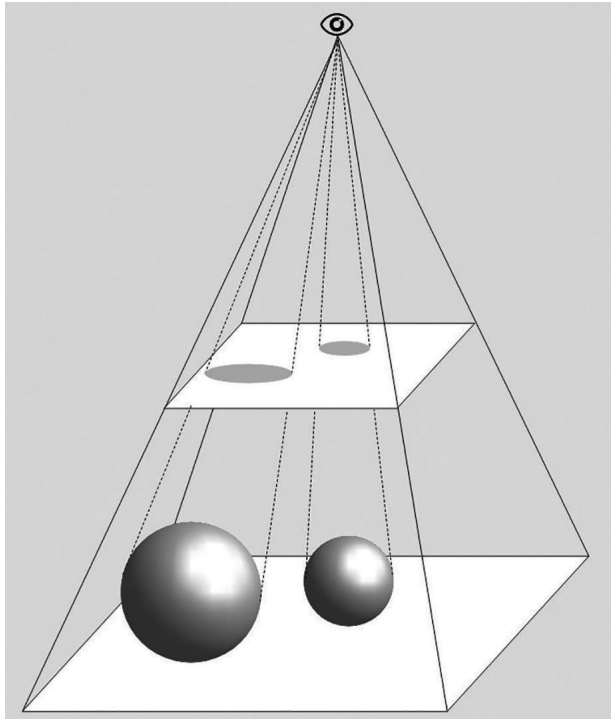

Figure 3.1
Constructing one among many possible Albertian Windows for a scene. A picture is a map of the light distribution at the window. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Distraction: *Primary vs secondary senses.* A picture of a model posing for King Arthur is in the primary sense a picture of the model and of the setting the model is in. The image is a picture of a person that at some moment in time sat in a room. This holds true, even if the picture is meant to, in a secondary sense, represent King Arthur, not the model. Pictures map the distribution of properties in space as available/projected/seen from a given viewpoint. The picture of a model posing, by its very nature, is a map of the visual properties of the setting in the primary sense.

Distraction: *Realistic pictures?* "What you say seems to apply only to photorealistic pictures, and even then, there are so many ways of generating photographs, each of which potentially distorts reality."[3] Yes, we talk only about photorealistic pictures.

Distraction: *Realistic pictures that do not represent real scenes.* Many photorealistic pictures, of course, do not represent a real-life scene. In these pictures, dragons can appear. There can be no maps of the Albertian Window of a scene containing a dragon because dragons do not exist. As I previously mentioned, we distinguish dragon pictures from pictures of dragons. But this is not crucial. The important thing is that the visual system feels sufficiently entitled to take a certain configuration as a map of an Albertian Window, and then proceeds to interpret it. This is a matter of degrees. A Mondrian painting, for example, could after all be taken as a map of an Albertian Window, thus as depicting some colored rectangles in a real scene, or as just a juxtaposition of colored quadrilaterals on the canvas.

Two important points to stress:

Point one: pictures are *not* maps of the scene they depict.

Consider: from the fact that pictures of an X are maps, it does not follow that they *are maps of X*.[4] A picture of my dog sitting on the carpet is a map of the visual array generated by light reflected, absorbed, emitted by the dog and the carpet, for a given viewpoint. This does not count as a map of the scene in which the dog and the carpet are situated, and in particular it cannot be used as a map of that scene.[5]

Point two: pictures are *not* pictures of the (generalized) Albertian Window.

This means that, *from the fact that pictures map the Albertian Window, it does not follow that they depict it*. Pictures do *not* depict the Albertian Window; they *map* it.

From these two points, a number of consequences can be drawn.

First of all, maps are *not good at supporting visual occlusion and perspectivality*. One element related to the dialectics of maps and pictures is the fact that pictures support occlusion and maps don't (or at least are not very good at it). Visual occlusion is a perspectival effect based on the relative location of two objects and a viewpoint. From where I sit, the window frame occludes part of the church apse across the street. Pictures are

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

excellent for the representation of occlusion because they are perspectival and thus bound to represent scenes that are full of occluders and occludees. This is why *topological entitlement* is not available for pictures. From the fact that the picture area representing the tourist's hand is in contact with the picture area representing the tip of the tower of Pisa, we are not entitled to infer that the tourist is touching the tip of the tower. Illusory contact is the price pictures have to pay for supporting occlusion.

On the other hand, maps can represent the colocation of two properties in space (see the discussion of superposed colors in chapter 2), but they are very limited in representing occlusion. A case in point is the representation of bridges over rivers (figure 3.2). If we translate into a formal map a portion of a map representing the bridge over the Oise River, we would have to represent the river as partly colocated with the bridge. But then, one would not know whether the bridge was above, below, or level with the river because we would lose occlusion. Occlusion is not well supported

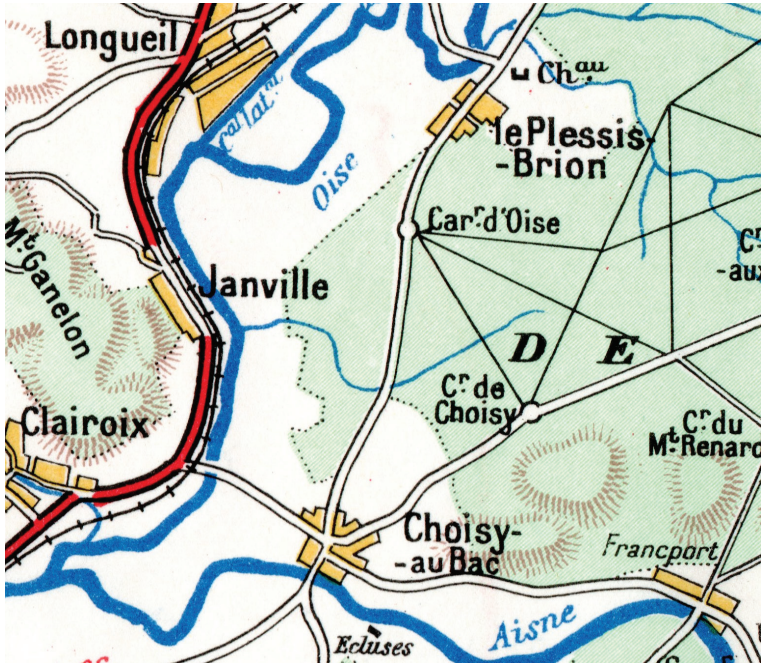

Figure 3.2
Rivers do not stop at bridges. Maps introduce pictorial conventions to account for visual occlusion. *Guides Bleues*. Paris, Hachette, 1952. Author's collection.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

by maps, except by convention or by locally resorting to picture-like depiction.[6] Incidentally, the fact that maps do not support occlusion aligns with the methodological constraint of same-dimensionality, which would also be violated by occlusion.

Given that pictures are maps of the visual array (point one), how are they apprehended when we look at them? The main informational job of pictures is to make a map of the visual array available for computational interpretation by the visual system. Given certain features in the map of the Albertian Window (e.g., systematic alignments), the visual system computes occlusions, which in turn are reliable indicators of relative distance from a viewpoint. Of course, the interpretation of aptly named *pictorial cues* is one of the most studied computational feats of vision.[7]

The Nature of the Albertian Window

The Albertian Window is a flimsy object. All that matters for the present account is that the Albertian Window, structurally, is a 2D surface. Once this is granted, I'm open to different interpretations of the Albertian Window's nature. If it is a physical entity with a precise location at a certain distance from the viewpoint of the image, then we can accept that there exist many Albertian Windows—each of them at a different distance from the viewpoint—and that the picture is a map of each of them. This highlights one of the advantages of the present account: the picture is both a map of many different Albertian Windows and a picture of a single scene. If the Albertian Window is some kind of abstract entity instead, then the picture is a map of an abstract entity.[8] Alternatively, if you do not want to consider an abstract entity as a full-blown territory, then pictures are not maps; they just behave in map-like fashion.

Maps and Structural Accounts of Images

To recap, pictures depict a scene and map the Albertian Window distribution of light from the scene. Pictures are not maps of the scenes they depict, and they are not pictures of the Albertian Window's light distribution. The account capitalizes on structural characterizations of depiction (Kulvicki 2014; Goodman 1968). Maps possess many of the structural features found in pictures. Maps and pictures are both syntactically sensitive. A map that

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

uses different colors for France and Italy has greater sensitivity than a map that uses the same color for both countries. Maps are semantically dense because they can use an unlimited color supply, and the supply is such that it is always possible to find an intermediate hue between any two hues to represent a property one seeks to represent. Maps, on the other hand, are, as a norm, not *transparent* in Kulvicki's sense. A picture of a picture of Sheena qualifies as a picture of Sheena, but a map of a map of Europe does not in general qualify as a map of Europe (although, in some circumstances, a map of a map of Europe may be used as a map of Europe, for instance, if the second map was requested by a typographer to map the color distribution of the first map; see figure 3.3).

The much-discussed Picture Principle, put forward by Fodor (2008) to distinguish maps from linguistic representations, states that if *a* is a picture of *b*, then parts of *a* are pictures of parts of *b*. The Picture Principle applies to maps as it is a straightforward consequence of the mereotopological constraint we claim for maps. The mereotopological constraint, on the other hand, does not apply to pictures, as the hand and Tower of Pisa meme shows that connectedness in the picture is not sufficient for actual connectedness in the scene. The Picture Principle has counterexamples in pictures; an item counts as a picture of Livio at the beach even if parts of the item represent the ocean background and not parts of Livio. If the Picture

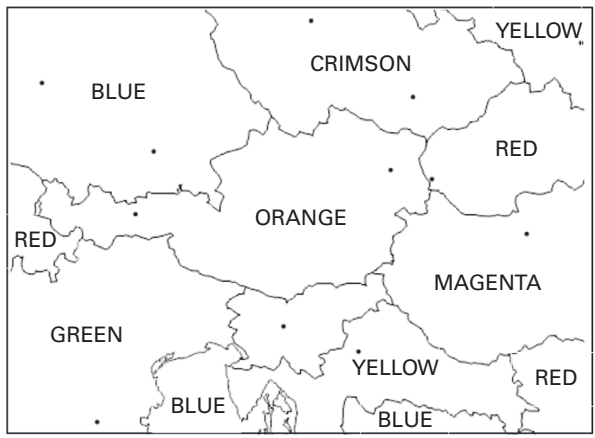

Figure 3.3
A map of a map of Europe. The typographer is given instructions about how to color another map. The first map will then be discarded. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Principle is a mark of analog representation at large rather than of pictoriality (Kulvicki 2015), then pictures turn out *not* to be analog representations. Maps, as a result, become prime examples of analog representation. The Livio/ocean background Picture Principle counterexample does not apply to maps. Maps map territories: if *a* is a map of territory *b*, then any part of *a* is a map of a part of *b*.

Pictures are also different from maps because pictures elicit recognition (Schier 1986), but they elicit recognition precisely by virtue of the fact that they are maps of the information perceptually needed for the eliciting. They could not elicit recognition were the visual information presented in something other than a map-like format. The case of hallucinatory "pictures," that is, generators of pictorial-like experiences that do not satisfy the systematicity constraint (Casati 2010), and the case of random dot stereograms proves indirectly the point: both cases fall in conceptual areas that partly overlap but do not coincide precisely with the conceptual area of pictures.

For a similar class, sound recordings are pictorial (Kulvicki 2006; Casati and Dokic 1994). They have a content/territory: an interview, for example, is the auditory content of a recording of an interview. If sound recordings are pictorial, the same duplicity of function applies to them. They are maps of the sound information, and they are pictures of the content. The recording of an interview is a map of the sound waves that are produced in the interview (from a given hearing point) and is an auditory picture of the interview.[9]

Pictures as Optical Maps and a Brief Discussion of Some Accounts of Pictures

I will engage in a little dialectics. The proposed account of pictures occupies a specific place in the general landscape of theories of depiction. At least one influential picture theory (Gibson 1978) bestows an important role to the optic array. But Gibson's picture theory fails to be a more general theory of pictures, and we are now in the position to understand why. Pictures are maps of the optic array, not pictures thereof. Conversely, the present account of pictures addresses at least some of the concerns about "sense data," "views," "aspects," "looks" (Noë 2006), "outline shape" (Hopkins 1995), and "occlusion shape" (Hyman 2006) theories of pictures. Consider

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

a picture, for example, that represents a round table from any position that is not above or below the center of the table. The picture is of a perfectly round table, but on the canvas we find (and measure!) an oval; the picture "resembles the look of the table" from a given viewpoint, but what are *looks* that they can be represented, and do we *see* oval looks? In hindsight, these concerns appear to be theoretical artifacts. At best, the theories embark on a rediscovery of the notion of visual angle, without clearly accommodating for it. At worst, they construe pictures as pictorial representations of these various items and generate an additional problem of making them "indirect" pictures of tables and chairs. But pictures do not depict visual angles (i.e., aspects, views, occlusions, and so on); they map those items. By mapping those items, pictures make eliciting visual recognition possible, and thus make depiction possible.

The account further eschews syncretism (Voltolini 2015), that is, the attempt to capitalize on seemingly competing accounts of pictures, in particular recognitional and structural accounts. The fact is, we do not need to reconcile different theories of pictures, as some of the items these theories describe are not pictures in the first place. The account proposed in this chapter further limits the need for dual theories of the content of pictures (Haugeland 1991; Kulvicki 2014), where pictures have both "bare-bones" content (e.g., the distribution of colors in various directions in a scene) *and* "fleshed-out" content (e.g., objects in various positions in 3D space having various colors). Given that pictures do not *depict* the distribution of light, they do not have bare-bones content; they still have, of course, fleshed-out content. It may be claimed, of course, that the issue here is only terminological or that maps are such that for them bare-bones content is identical to fleshed-out content.[10]

Another consequence of the present account is that mental "images" are better construed as maps, not pictures. This claim addresses the problem (Kulvicki 2014) of the regress of mental image interpretation. If mental images are construed as pictures and if pictures are not construed structurally, but rather as perception-dependent items, mental images would require for their interpretation a homunculus endowed with perceptual capacities, leading to a regress. But if mental images are construed as map-like, then all they are requested to do is to make information available in a map-like format, whichever the use perceptual and other systems can do of the thus formatted pieces of information.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Maps, Pictures, and the Picturification of Maps

As noted, pictures and maps are visually germane to each other. When you look at a map, you can take it for a picture because it looks like a picture or it can be made to work like a picture; in turn, these two features depend on the fact that maps and pictures share a number of syntactic and semantic features. Now, taking a map as a picture is an illegitimate interpretive move. A territory's visual features are not extractible from the conventionalized visual features of the map (houses do not have pink roofs, at least not all of them; roads are generally not yellow or red all over; and so on) unless the visual features of the map come so close to, let's say, a satellite image's visual features that we are unable to make any clear difference between the map and the satellite image. Moreover, the visual features depicted in a photograph may hide important similarities between parts of the territory (e.g., the sea in a photograph can be of many different colors because of differences in incident light).

Some maps can look, or can be visually manipulated to look, more like pictures of the item they map than just maps. This is the case for nano "images" that are, in fact, cartographies of a certain small physical area produced by the successive recordings by a sensor. In these nano "images," we find visual artifacts such as nonexistent bumps and hollows created by readings of positive and negative potentials and shadow-like appearances generated by the sensor.[11] This can also be the case with physical maps that introduce shadows to enhance the perception of mountainous reliefs. In maps, shadows are conventionalized as coming from a northwest Sun position that takes advantage of the visual system's interpretation of concavities and convexities, presupposing light from above. As a rule of thumb, a clear symptom of the picturification of maps is the addition of shadows to enhance the 3D appearance of the represented territory. In these cases, the map designer takes advantage of the fact that pictures are Albertian maps, which makes maps easy to mistake for pictures. Picturification happens in written language and in diagrams, too. Again, the tendency to take these items as pictures should be resisted.

Reading Aerial Photographs: The Southerner's Advantage

Those who think that maps can be replaced by aerial imaging should be aware of some biases that plague pictorial perception of concavities and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

convexities. Aerial photographs of some of the most mountainous areas of the northern hemisphere are subject to a concavity/convexity reversal illusion, due to three facts: sunlight comes from the south (except above the Polar Circle in some parts of the year), the visual system is bound to interpret concavities and convexities as if light always comes from above, and north coincides with above (spatially) in the pictures. This fact is particularly sensible now that aerial photographs are ubiquitous.

Consider the following satellite image (figure 3.4).

In a desert landscape, mountains suddenly rise with sharp crests, and layers of rock upon layers of rock. It is a strange landscape, come to think of it. An ambitious road has been built on the highest crest; the construction seems incomplete toward the right of the picture. The deep valleys would host rivers, but where do these rivers flow when the valleys meet the plain?

A geographer would say that this is an upside-down world of sorts. And indeed it is, but only from a cognitive point of view. And geographers, of course, know better, as the image depicts one of the most prominent features of Earth. This is actually a north-up picture of the South Rim of the Grand Canyon, Arizona, west of Grand Canyon Village. The "road" is the

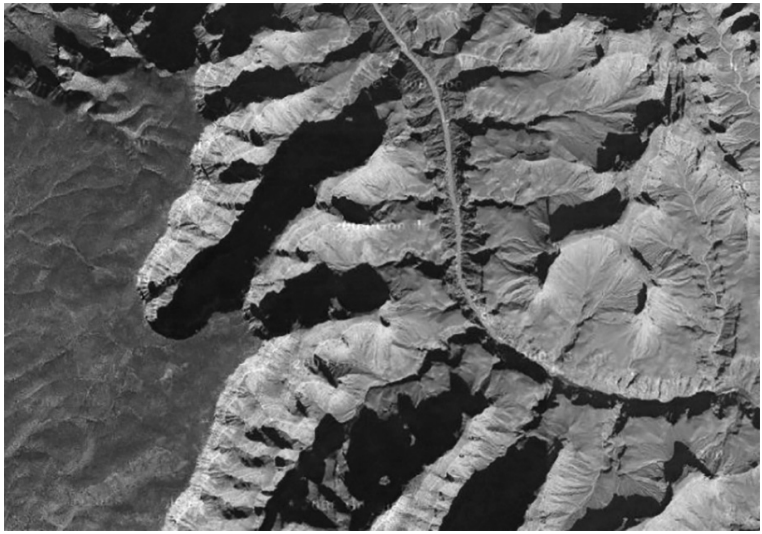

Figure 3.4
Steep geological features. Retrieved from maps.google.com, November 2011.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Colorado River, and what you see as crests are actually the bottoms of valleys. No water flows into the plain from the valleys because the Coconino Plateau is the highest feature in the picture.

Now that you know what is in the image, can you mentally "reverse" the picture, so as to see elevations in the right order, that is, mountaintops closest to your eye? This is not easy. Research in vision science has shown that the visual system makes extensive use of an implicit rule in retrieving shape from shading. The rule states that light must be assumed to come "from above" (Kleffner and Ramachandran 1992; Mamassian and Goutcher 2001; Casati and Cavanagh 2019). The rule is consistent with ecology—in the natural environment, light overwhelmingly comes from above—and thus has a strong evolutionary plausibility. "From above" appears to mean literally "from above" when one looks at a photograph on the screen of a laptop in standard viewing conditions, but it can also mean "from the best approximation to from above" (e.g., when one looks at a picture that lies flat on a table). If unconstrained, the interpretation of shading to retrieve shape is underdetermined. A shaded area could be induced by a protrusion that is situated immediately below it and blocks light from below, or by a protrusion that is situated immediately above it and blocks light from above. If light is assumed to come from above only, the first interpretation is eliminated. The fact that we typically see pictures such as figure 3.4 as a pattern of convexities (mountains) rather than a pattern of concavities (a canyon and its tributaries), together with the fact that even if we are told, it is hard to revert to the correct interpretation, is taken as an indication that the rule "light from above" is deeply ingrained in our visual system.

This fact, however, is an indication of difficulties faced by readers of aerial photographs. Most[12] places inhabited by Northerners see the midday Sun in the south. Aerial photographs taken at that time of the day (to minimize the effect of cast shadows) and that are oriented so that their north coincides with the top of the picture, will display a pattern of "light from below" and, accordingly, will incur the risk of a concavity/convexity inversion. The places inhabited by those who live in the southern hemisphere will not be subject to this effect as they (mostly) see the midday Sun in the north. The convention of representing north in the top part of the picture thus induces an asymmetry, that results in a Southerner's advantage when it comes to aerial photographs.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Maps typically correct for the asymmetry in representing shadows as cast from a northwest light that is impossible in the northern hemisphere. But what remediation can be proposed for aerial photographs?

There is a way to put the 3D unfolding of the picture in the right order, so that the highest features are closest and the bottom features farthest from our eye. You should simply turn the picture upside down. You can try this with figure 3.4. The effect may not be immediate, but when it is in place, it is robust in the sense that the concavities and convexities stabilize and it is near impossible to see again the depth of the Inner Canyon and the Colorado as a raised crest with a road.

Thus, the remediation consists in learning to use north-up aerial photographs of areas in the northern hemisphere by turning them upside-down. This is something we can do, as the alignment of north with the top in a map is the result of an unlearnable convention. We cannot, on the other hand, train our brains to unlearn the "light from above" assumption.

However, this remediation may prove difficult if we work with the screen-anchored visual representations—as millions of people do today when using the popular geographic websites. It is not always possible to rotate the viewing screen, and taking a screenshot, pasting it into an image editor and turning it upside down requires too many clicks. May we suggest that web providers of satellite images insert a simple button "Turn 180°" in the interface? This would also have the indirect advantage of loosening the cultural association between north and top of the picture.

What Came First, Maps or Pictures?

Are maps pictures whose perspectivality has been abstracted away? We do not directly access mental maps the way we access visual images of a scene. And before the invention of human flight, we never had consistent or reliable access to a viewpoint from high up. How did mappings happen, then? Was it by laying pictures of panoramas on a table? Was it by realizing we could turn pictures of panoramas into maps, with little adjustments? Was it by imaginative play? Mapmaking is different from picture making because the Albertian Window for a territory was not accessible before the invention of flight. People were able to reconstruct images of the Earth as from above, but these images remained reconstructions.

This reconstructing is what astronomer Camille Flammarion purported to do in his 1880 made-up picture of the Naples area (Flammarion 1880),

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

which he evidently derived from a map (figure 3.5); see (Casati 2003a). Cartographers needed (and still need) to measure the territory and to report their measurements onto paper (or digital facsimiles). Moreover, the cognitive practices linked to map understanding are different from uses linked to picture understanding. As we have seen, hypergenerativity does not require recognition. Therefore, I lean toward saying that maps are not derived historically from pictures. But this is a historical question.[13]

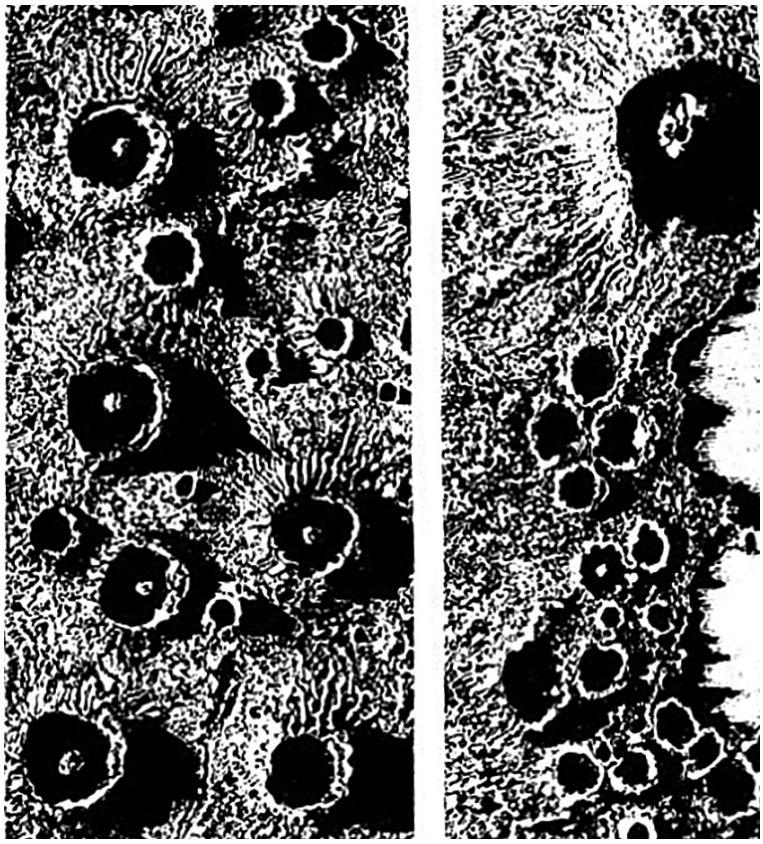

Figure 3.5
(Left) The moon. (Right) The Mount Vesuvius and Naples area. A reconstructed aerial view. Camille Flammarion, *Astronomie populaire*, p. 166. Author's collection. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Intention Detection, Artifact Categorization, and Hypergenerativity

According to accounts of artifactuality (Bloom 1998; Pignocchi 2010; 2012), intention detection is active each time we construe an object as an artifact (i.e., as a human-made item). Intention detection may actually be responsible for the mental labeling of the object as an artifact. Intention detection can also misfire in the case of natural signs (i.e., shadows, traces that are left unintentionally). Its contribution to the mechanics of artifact understanding is generic and not specific to maps. We can see that a certain object is a map, in the sense that we understand that it has been designed as a map. But given what we know about hypergenerativity, we also see that the existence of referential relations between parts of the vehicles and parts of the territory is not a consequence of intention but a consequence of map's spatial structure. This claim also holds for the part of the story that concerns the choice of a projection method. Once more, the intentionality factor is overarching but harmless. There is an intention behind the choice of a projection method, but once a projection method is chosen, there will be consequences, and these depend on the explicit use of the vehicle's spatial structure, not on the intention.[14]

Maps Not Usually Considered as Maps

Having dealt with the issue of the picture/map border, let us deal with some examples of items that the present account reads as maps, and of items that lie at the boundary.

Eliciting Mappings: Tables of Contents versus Indexes

It may take some ingenuity to work out mappiness in certain cases (and here a detour through formal maps that can help). Consider the opposition between a *table of contents* and an *index*. Both list the printed materials within a book (their common territory) by specifying the location of that content. Book pages are numbered progressively, and both the table of contents and the index report the relevant "address" for the sought-out information. Now, tables of contents, according to the present account, are maps, whereas indexes are not. Though both are mappings (from list of content descriptions or names to locations where the content is located), only the table of contents is spatially organized in a map-like way. The

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

part of the table that refers to chapter 2 is between the parts that refer to chapter 1 and to chapter 3. The part of the table that refers to chapter 2 is contiguous to the part that refers to chapter 3, and so on. As a result, the same-dimensionality constraint is respected, and we are able to locate ourselves in the table of contents (e.g., "I just finished reading chapter 2, let me mark it in the ToC so now I can see what I am supposed to read next.") The organization of the index, on the other hand, is alphabetic, referring to a mental representation of alphabetic order—without which it would be difficult to use an index at all.[15] Thus, a line in the index referring to a piece of content on page 56 may be next to a line referring to a piece of content on page 83. No spatial organization is requested from the index, except the abstract mapping to the alphabet. Observe that tables of contents do not need numerals to refer to book parts. By looking at a table of contents we can navigate the territory—the book—although this navigation may take more time than if page numbers are available. On the other hand, page numbers remain essential to indexes.

Granular Mappiness: Tables of Content Once More

The semantics proposed by Camp (2007) includes the concept of a *degree of resolution*. Degree of resolution is missing from our earlier account (Casati and Varzi 1999) and is a welcome improvement. The truth of a map is relative to a degree of resolution. Some maps, like seating charts in airplanes, cannot show whether you will sit askew or lean back, or whether the seat will be too big or too small for you. Airplane seating charts only show the location of a seat relative to other seats, and they preserve the topological relations between seats so that you can locate and assign seats unambiguously. We can improve on our original map semantics by introducing a similar clause. Some maps have a lower degree of resolution than other maps of the same territory. This means that in the territory there are "atomic regions" whose internal mereotopological structure would not be visible to the map. Similarly, tables of contents in books are mappy in a "granular" or "quantized" way. Accordingly, we can only draw *limited spatial inferences* from a representational vehicle's organization. Consider the noun phrase "Chapter Three" that we can find on the table of contents. Its left half, "Chapte," does not refer to a part of the book that is adjacent to the referent of its right half, "r Three." Thus, certain areas in the table of contents can be singled out as

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

they are occupied by a certain linguistic expression (typically this is a singular term, or possibly a quotation of the corresponding items in the book).

Chapter One
Chapter Two
Chapter Three

Each boxed area represented above is a granule or quantum of spatial reference. The space in the table of contents is discretized, and each area refers to an area in the book-territory, without a subpart of the vehicle area having any particular semantic significance. By looking at a table of contents, you know that certain contents will be next to, come before, or are included in other contents.[16]

Things That Do Not Look like Maps but Are Maps: Brain Scans and Nano "Images"

Certain items are commonly called "images" but should rather be categorized as maps. Brain imagery items (e.g., obtained through fMRI) are prime examples of the picturification of maps. They are displayed in books and journals as pictures of the brain, with no clear marks indicating that they should be read as maps. In fact, some of the visual features of brain imagery make us spontaneously apprehend them as pictures. However, there is no *visual* aspect of the brain actually represented by the colors in brain scans. The colors map areas that have certain distinctive properties such as levels of blood sugar concentration relative to a set baseline. The levels are represented by colors whose choice is dictated by visual convenience (Roskies 2007; Legrenzi and Umiltà 2011). Representations of nanoscopic items are obtained by a number of processes that probe details of the physical structure below the visual wavelength threshold (Chaput 2019). The results of the probes are then spatially organized. Some visual accidents give viewers the impression of shadows, and this in turn contributes to the picture-like appearance of nano-images. These representations ought to be renamed "brain maps" and "nano-cartographies," respectively.

In both the case of brain images and nano-images, the representations are the result of *recordings* (Kulvicki 2014) produced by automatic and mechanical reactions of a device to environmental states. Geographical maps have historically been produced by non-recording devices (measurements, drawings)

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

operated by trained staff, or if they did involve recordings (e.g., by relying on aerial photographs), trained staff would produce maps indirectly (e.g., by analyzing and comparing different photographs). If certain items are called "images" or "pictures" because they look like pictures, then false coloring, the abolition of shadows, and the use of quantized scales (instead of smooth gradations) help keep map picturification at bay. Kulvicki (2010) suggests rendering brain "pictures" more cartoonish to prevent their being taken for photographs or for representing visual aspects of the brain.

A closer conceptual equivalent to brain scans and nanoscopic representations are geological maps when they are produced by displaying the results of measurements made through probes.

More recent techniques allow automatic map extraction from recordings. If mechanized, automatic pictures (photographs) are also a (rare) instance of an automatic map—once more, a map of the Albertian Window—automatic maps of other territories, in particular of territories typically mapped by geographic maps (i.e., maps of Sardinia or maps of Central Berlin), are technologically challenging. Some recent mapping developments use crowdsourcing, i.e., distributed sampling by users to automatically and mechanically generate maps (e.g., real time car traffic maps). An image of the paths left by tracking systems (AIS) on vessels at sea will inevitably make the coastline emerge as a by-product of the fact that the coast limits the boat's movement. But in general, it is difficult to directly generate a geographic map of a scene from a picture of that scene (it is easier to generate a map from multiple pictures), and indeed there exist picture-based locational algorithms (Chiabrando et al. 2017). There are not many direct mapping devices that work like cameras: click on the shutter, get the representation. Cameras generate pictures by recording the map of the optical array, and then this map is interpreted by the visual system that reconstructs the scene based on likely occlusions and other visual cues. Automatic map-generating devices ought to measure distances in order to recover and make available spatial relations between objects. There do exist non-centered 3D captors of 3D territories such as particle trackers in particle colliders or wind sensors that are distributed over a territory. There are other types of captors (i.e., medical ultrasounds, radar images) that are not organized as sheets but that deliver the information in a map format.[17] Thus, unlike photographs, maps are not defined by the method through which they are generated. FMRI, nano-maps, and GPS algorithm-generated cartographies are examples that come close to mechanically generated maps,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

without matching the causal immediacy of photographs. This is in general to be expected because maps are the result of abstractions.

Kulvicki (2020) claims that whenever you look at a map, you take it to be a factive representation, as opposed to what happens with nonphotographic pictures. Kulvicki's argument is an empirical claim that we can endorse for the time being. Its explanation requires some work because maps are not mechanically generated the way photographs are.[18]

"Good" Picturification: Visual Reconstructions

Some representations are generated in the way maps are typically generated: by measuring features of a territory, displaying them on a map, and then using those very measures to produce realistic visuals of the item represented. This happens when the represented item has, in principle, visually accessible features. We can send probes in caves, on the far side of the Moon, or at the bottom of the sea and generate visuals based on the rock samples they collect. Making these visuals photorealistic has potential advantages over keeping them map-like. For instance, cave explorers (Moyes 2002) will find it convenient to use more than just a map of the territory.

Cross-Sections and Axonometries

Earlier, we claimed that cross-sections (figures 3.6 and 3.7) such as those of geological piles, architectural buildings, vessels, and geographic surfaces may count as maps. However, these cross-sections do not have a pictorial correspondence due to the fact that there is no viewpoint from which the pile or entire building is visible as it is in the map. Rather, each cross-section would be, in a sense, the prolongation of the sequence of Albertian Windows into an object.

A 3D architectural model (possibly carrying inscriptions of quotes) would preserve all the spatial properties of the represented building (modulo similarity), *and* it would be visually convincing (see chapter 5). But a 2D representation must, by necessity, sacrifice some properties in the interplay between accuracy and visual poignancy, between informational richness and mental computability. If one wanted to preserve visual shortening, then metric features and parallelism will go. On the other hand, preserving metric features creates very partial «views» (i.e., side views, top views, and so on), each of which only offers a necessarily limited vision of the object.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

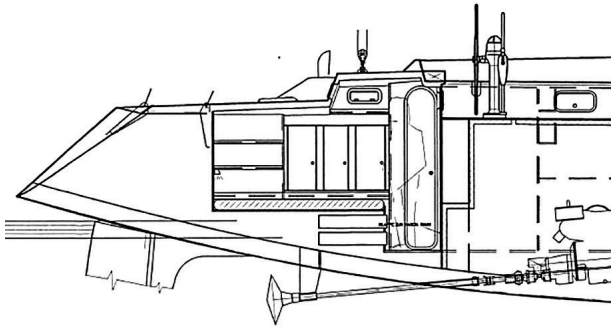

Figure 3.6
A cross-section of the stern area of a vessel maps the light (or contour) distribution as if the object was cut. Author's collection. Image credit: RC.

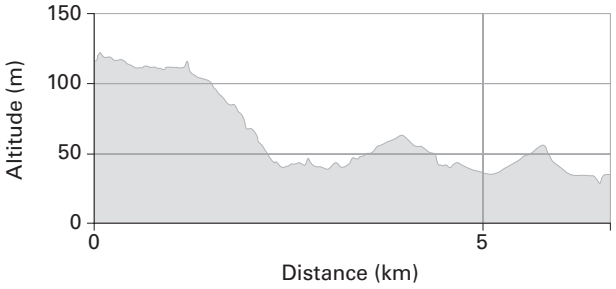

Figure 3.7
A south-north elevation profile of the Paris area (Villejuif to Notre Dame). Image credit: RC.

Axonometry (figure 3.8) is an interesting compromise (Tüscher 2006) that allows for effective and satisfying visual inspection, while at the same time allowing for the preservation of relative scale.

Discussions of axonometry (Hagen 1986) make use of the notion of a "viewpoint located at infinity." This locution is meant to capture the idea that in axonometries no visual shortening is involved. Shortening is indeed a consequence of a viewpoint located at a finite distance from the object. However, the mathematical concept of location at infinity is metaphorical. In the context of an ecological description, there is no room to make a physical interpretation of the notion of a viewpoint located at infinity. No such point could physically exist or be usable for depicting an object. The absence of shortening in axonometries has a different physical grounding,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

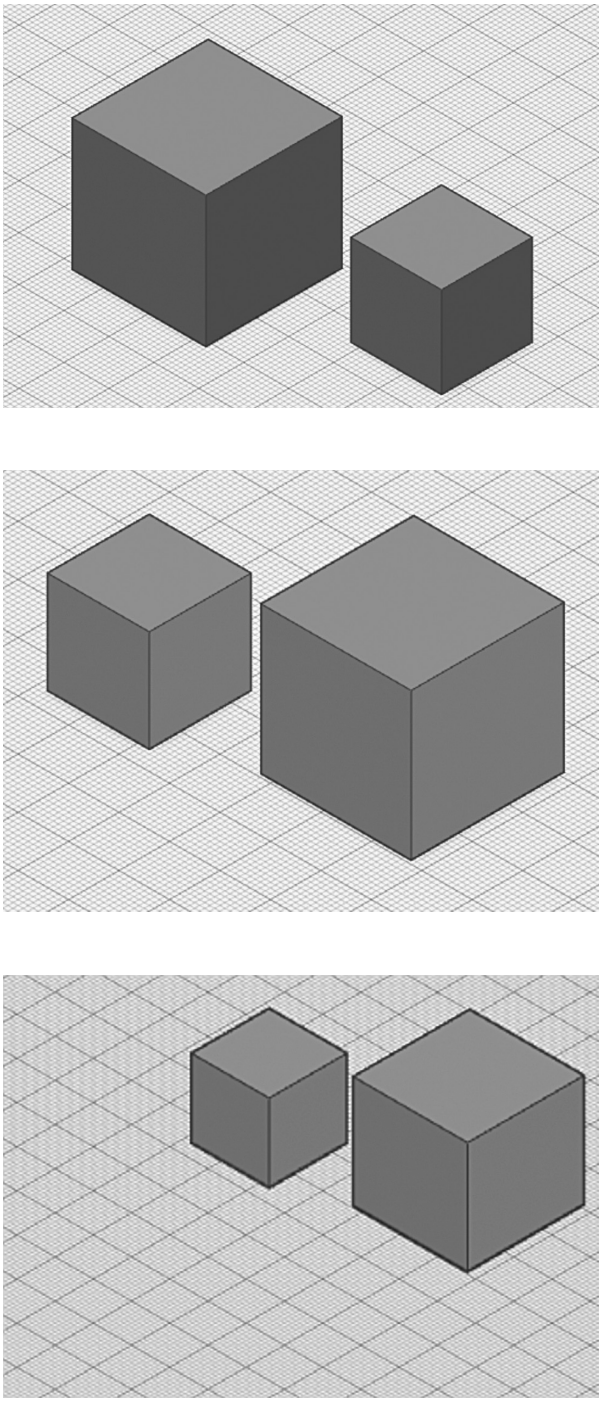

Figure 3.8
(Top, middle) The same scene from two different directions of view. (Middle, bottom) Moving along the direction of view (zooming out) does not produce foreshortening. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

that is, the fact that projection rays (or light rays) are parallel to each other. The shared feature of parallel rays is their *direction*. I propose to consider *direction of sight* as the defining feature of axonometric representations. Axonometric representations abstract from viewpoint, and only rely on a direction of sight (figure 3.8). Their reliance on a direction of sight is important as it marks a difference from even more abstract representations, such as mathematical models of the represented objects (e.g., a set of ordered triples of numbers, with each triple characterizing a point of the object using the x, y, z, axes). Direction of sight allows to account for *visibility*. In axonometries, a minor adjustment of the proposed notion concerns the fact that the direction of sight is *oriented*. It makes a visual difference whether your axonometry shows the top side and two vertical sides of a cube, or whether it shows the bottom side and two sides of the same cube (Casati and Giardino, 2013).

The next step in our argument is that at a certain distance, objects are seen as good approximations of axonometries (e.g., optically zoomed images, satellite images). The exact determination of this distance is a matter for empirical research, but we can provisionally consider[19] a ratio of 1 to 20 between the angular size of the object and its distance from the observer. This means that if an object of angular size of one unit is at 20 units of distance from the observer, parallel lines on the object will not be represented by converging lines in the image but rather by lines that will be sufficiently close to parallel so as to be seen as parallel.

However, when in the visual field many objects are simultaneously present, their overall angular size will be, in general, a much larger fraction of their distance from the observer. This entails that the local (quasi)-axonometries will not survive across objects that are far apart in the image. Each object will be seen axonometrically, but the respective images will not merge in an overall axonometric perceptual representation. For each part of the object a unique *direction of sight* is specified; and each object has its own direction of sight that is different from that of any other object. This further entails that precise visual cross-object comparisons will be hard, in the same way they are hard when considering a photograph of two objects that are quite apart in a given scene. Even if we cut the area corresponding to one of the two objects and bring it closer to the area corresponding to the other object, shortening and perspective will give us a hard time in making the comparison.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Ambient axonometries remove this difficulty. By completely abstracting from the viewpoint, axonometries force a unique direction of sight for *all* objects in the represented scene. Consider a typical axonometric representation of a city (figure 3.9). The sacrifice of perspectival verisimilitude (the abandonment of shortening) makes it possible to perform the cut-move-and-compare operation *across the whole surface of the representational vehicle* (figure 3.10).

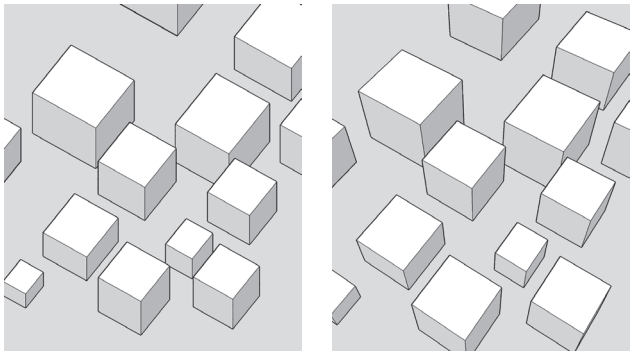

Figure 3.9
The same simplified cityscape represented in an axonometry (left-hand pane) and in perspective (right). Axonometry preserves parallelism and ratio of the sides of the boxes. Perspective is visually more convincing, more immersive, and implicitly hints at the viewpoint. Image credit: RC.

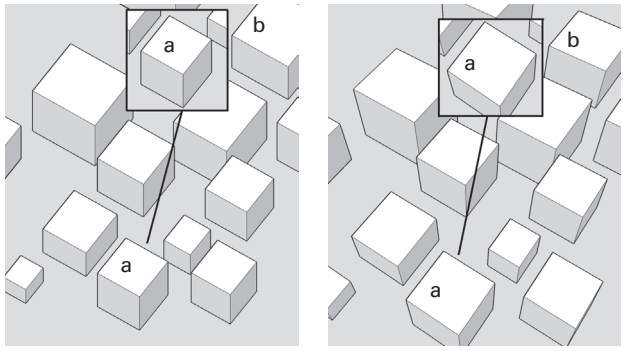

Figure 3.10
Visual comparison of elements of the scene (here a partial view copied and pasted next to the element to be compared) reveals axonometry's advantage. In the axonometry (left pane), we realize that b is much larger than a, and that the difference is more salient than it is in the perspectival rendering. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

The *specific* representational advantage of axonometry is thus *long-distance comparison*. It is an advantage that is related to, but stronger than, other classically described advantages, such as preservation of parallelism, and it is an advantage that axonometries share with maps.

There is an axonometric effect in zoomed pictures, and in satellite photography. The latter shows that axonometries also come to play in the map ballpark: an axonometry is a picturified map of a territory. The distribution of properties is map-like, whereas the coding of elevations is picture-like. The advantage created by the picturification of elevations is their immediate readability; this is only enhanced by the possibility of long-distance comparisons (figure 3.10).

Celestial Maps Are Not Maps of Stars

Celestial maps (figure 3.11) appear to violate the same-dimensionality requirement. The stars they represent are at wildly different distances from Earth. But celestial maps are actually maps of the visual information (of the visual angles subtended by pairs of stars or by manifolds of stars.)

Celestial maps neatly exemplify a need to block the inference from being a visual map to being a map of the thing represented. Celestial maps are not maps of stars, but they map a very peculiar vault-shaped Albertian Window above our heads. They are maps of the visual distribution of starlight coming to our viewpoint on Earth. They qualify as a *picture* of the stars. You cannot use a star map to plan an energy-efficient voyage that touches both Rigel and Aldebaran, as the map does not specify which star is closer to Earth. With Rigel at 860 and Aldebaran at sixty-six light years away from Earth, such specifications would make a significant difference. Historically, celestial representations were conceptualized as maps in periods where the sky was thought to be a 2D vault. Had the sky been a 2D vault, spatial relations between stars would be mapped by spatial relations in the representation.[20]

Non-Maps: Diagrams

There is a lot to say about diagrams, and of course a lot has been said about them, in particular in the Peircean semiotic tradition.[21] Here I am only concerned with the way they contrast with maps. Diagrams use spatial vehicles to represent a certain set of facts, objects, and the relations between them. They must be M2-read—a complex interpretive process that incidentally

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

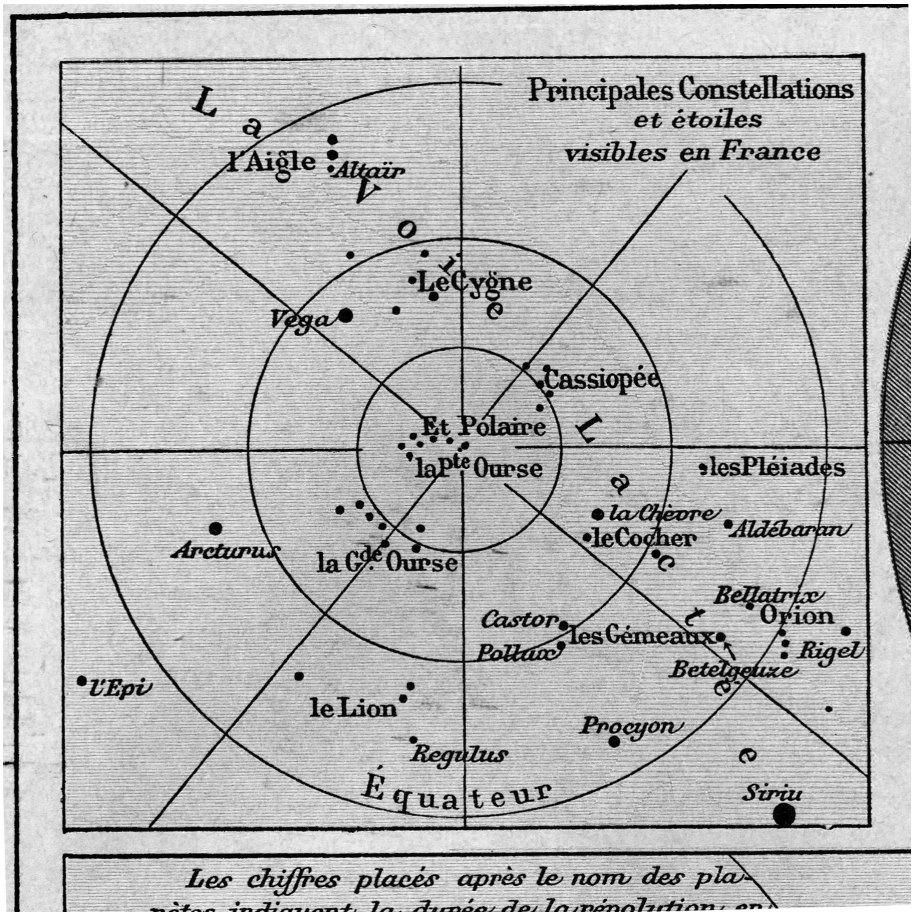

Figure 3.11
A celestial figure of the northern sky. Image source: *Atlas classique*. Paris: Hachette, 1931. Author's collection. Image credit: RC.

requires inhibiting looking at them as maps even if they represent a spatial territory. After a while, diagrams are maybe accessed in M1 mode, but no longer as if they were maps. The visual shortcuts that are used in diagramming are not map-specific visual shortcuts. Diagrams, such as histograms, are not maps (though they may well qualify for the title of mappings). The contrast with histograms is pretty straightforward in most cases. The two-bar histogram representing the figures of the dog and cat populations in Paris does not represent *where* dogs and cats are in Paris, or where they

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

are relative to one another. A portion in the left bar does not represent any specific dog or group of dogs, nor does it purport to represent any of their individual locations. No locational inferences can be drawn from the histogram. It does not matter whether the dog bar is on the left or on the right. Similarly, the spatial relation between the data column on the left representing the Greek population in 2015 and the data column on the right representing French population in 2015 is not meant to reflect a spatial relation between the two populations. Adding the data column for Uruguay on the further right proves this point. Even a more "geographic" diagrammatic representation of a set of spatial entities, such as rivers, can fail to be a map (figure 3.12).

You can even draw spatial inferences from diagrams, without turning diagrams into maps. A histogram representing the GDP of Argentina and of Switzerland will entitle you to infer that there will be a prevalent migratory pattern from Argentina to Switzerland and not conversely. Still the histogram under discussion is not a map. There is more than usability for drawing spatial inferences that characterizes maps.[22] Thus, an analog representation of a spatial territory, using space as a representational vehicle, and from which spatial inferences about that territory can be drawn, may not be a map.

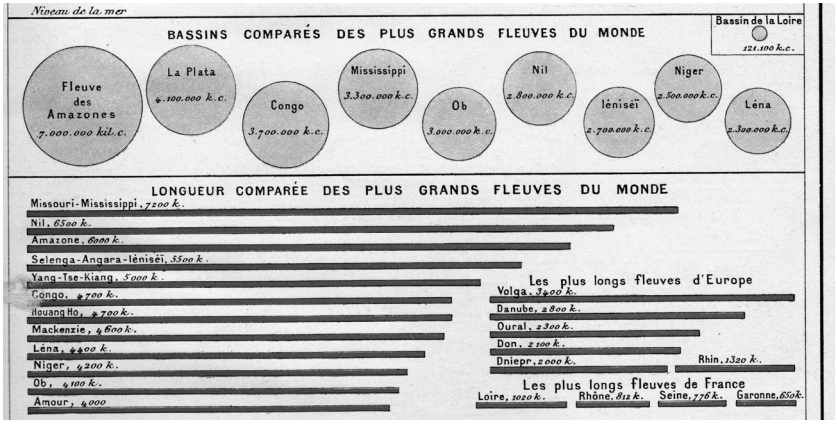

Figure 3.12
Rivers are geographic entities, and arguably the representation of each of them by a line counts as a map. But the diagrammatic juxtaposition of linear maps of rivers is not a map of the rivers of the world. Image source: *Atlas classique*. Paris: Hachette, 1931. Author's collection. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Conclusions

I made a number of claims in this chapter: pictures are maps, in a special sense. They are not maps of what they represent, and they do not represent what they map, but they are maps nevertheless. We have explored commonalities and differences of maps and pictures and the map-like nature of certain types of visual artifacts, including diagrams. Some of the next chapters will explore other maps and map-like structures. Before going into that, we'll take a look at the working of maps in navigational contexts.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

4 Maps, Navigation, and Situation: Red Dots, Green Dots

As we have seen, maps can be used in situated and non-situated ways. One can see on a map that the distance between Lisbon and Paris is shorter than the distance between Lisbon and Tokyo, or that Madagascar is an island (i.e., it is disconnected from continental Africa).[1] However, one may also need to use a map for navigation (figure 4.1). In this case, the map needs to be situated. But what is situatedness? *Situating* a map consists in literally having some syntactic element of the map *physically coincide* with the portion of the world that element is supposed to represent. Thus, not only is the situation different from representation. A map can represent a territory it is not situated in (e.g., a map of Tasmania hanging in my office in Paris). It also appears that purely representational relations are not sufficient for navigation. What follows examines this claim and proposes a weak, epistemic reading of situatedness.

According to a strong view, maps' *representational capabilities* depend on the holding of some nonrepresentational relations between the map and the world. In a world that displays massive reduplication in spatial patterns (Strawson 1959), we may have no way to decide whether at the dorms this is the map of my room or the map of your room. A map that represents rue Saint Jacques intersecting the axis of the Notre Dame Cathedral could be a map of Paris or of Montreal. Whether it is a map of one or of the other depends on extrinsic relations. For example, it has been drawn after measuring places in Paris, not in Montreal, so it is a map of Paris, not of Montreal. According to a weaker view, some map-to-world relations cannot be purely representational if the map is to function for *navigation*. There have to be physical, nonrepresentational relations that connect the map to the region it represents. *Not* in order to *represent* that region, but in order for the map to be *used* for navigation in that region. We explore this weaker option in what follows.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 4.1
How to go from here to there. Paris, July 2018. Image credit: RC.

Navigation and Situation: Red Dots

Consider a map with a Red Dot on it, to be interpreted as asserting "you are here," as per the map legend (or maybe as asserting "this map is here," "the map you are looking at is here." The choice of the precise formulation could be a minor Distraction. Another Distraction is the color of the dot. On static maps, the dot used to be red. Nowadays, the dot is predominantly blue on GPS-based navigation devices. A notational variant is the map-marker symbol $\heartsuit.)^2$ To be used for navigational purposes, a Red Dot map must be either a static map *nailed* to the ground or a dynamic map (e.g., GPS related) continuously updating the position of the Red Dot as the device moves

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

(i.e., by shifting the map on the screen). If a static map of Navona Square in Rome was not nailed to the ground in Navona Square and you could carry it to Naples, you would be misled by the "you are here" Red Dot that now, in Naples, suggests that you are in Rome where you, in fact, are not. A dynamic map makes the dot portable, but it must have it move relative to the map as you move on the ground to keep representation and situation in sync. In both static and dynamic cases, the Red Dot can only work for navigation because *it* is *in* the place it represents.

The nonrepresentational relation (figure 4.2) the map bears to the world is, in this case, plain *identity* or *colocation*. The Red Dot coincides with the object it represents or the place it occupies. From now on, I'll capitalize "Red Dot" when it comes to its navigational function on a map.

Navigation and Situation: Green Dots

Nonnavigational uses of the map can abstract from nonrepresentational relations such as identity. John, on his visit to Rome, wears a geolocation device that broadcasts his position. From Paris, Paula can track John's position on the ground in Rome. She observes a Green Dot corresponding to

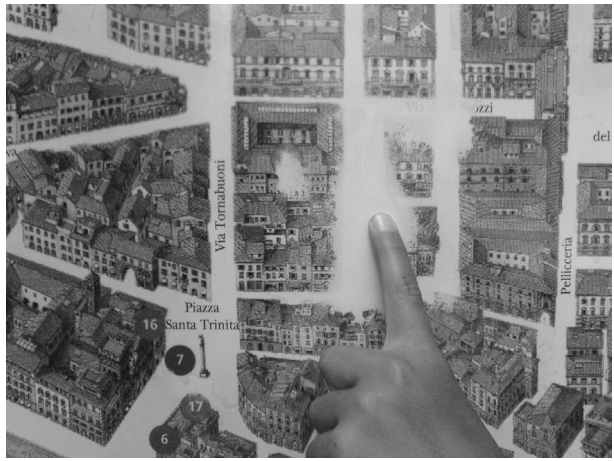

Figure 4.2
The physicality of our cognitive engagement with the Red Dot is apparent by the fact that the dot is oftentimes erased by people putting their fingers on it. Florence, July 2014. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

him on *her* own map of Rome that she is using in Paris. The Green Dot refers to places in John's environment, but it is not identical to or coincident with any of the places in John's environment. Conversely, Paula cannot use the map to navigate *her* environment in Paris. Sure enough, as we said earlier, she can read off a map many facts even if the map is misplaced. (Navona Square is correctly represented as included in Rome on the map, even if the map is in Paris.) The fact that the map is dynamic and constantly updating does not make it a navigation device. Navigation is possible only if the map is situated.

The situatedness claim can be modulated in a number of ways. Can John use the non-situated Parisian Green Dot for his whereabouts in Rome—for instance, by asking Paula what to do and taking directions from her remotely? Yes. But taking directions is not navigating; Paula is the one who does the navigational computations. What if John can visually access the map of Rome (with his location marked on it using the "you are here" Green Dot) that Paula uses in Paris, while still being at some distance from the map, possibly looking at it through a webcam or videos sent of the map? In that case, the Green Dot's representation (in Rome) of a Green Dot (in Paris) is arguably the functional equivalent of a Red Dot in Rome. (Let us assume, for simplicity's sake, that there is no noise in the iteration and no temporal delay.) The claim would be that the Green Dot John sees on the display of his computer in Rome—which is just a representation of the Green Dot Paula sees on her computer in Paris—works like a situated Green Dot, and hence is a Red Dot. It is in Rome, and it tracks John's position in Rome. But the representation (or a copy) of a situated representation is not thereby a representation with the same situation as that of the represented representation. The point I wish to underscore is that John must *know* that that is a public representation of his location in Rome.

Distantly Accessed Maps: Epistemic Segregation

One can imagine an even more extreme case here. The display of the map (the one Paula is staring at in Paris) with John's location gets projected on the Moon, and John can see and inspect it from where he is. One may be tempted to say that in such a case, John *can* use the map for navigating in Rome after all. But on second thought, one should convince oneself that John needs a lot of supplementary evidence to secure the knowledge that the display on the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Moon is a map of *his* current location. Consider: the Moon is visually accessible from a whole hemisphere of the Earth, which includes John's current location. If someone saw the display on the Moon from Greece, for example, one would not be *entitled* to use it *there* as a guidance for navigation (figure 4.3).

What is missing here? In most ordinary cases, the proximity of the map to John's actual location provides an epistemic warrant for its use. The label "you are here" or "this map is here" is meant to work because the label is

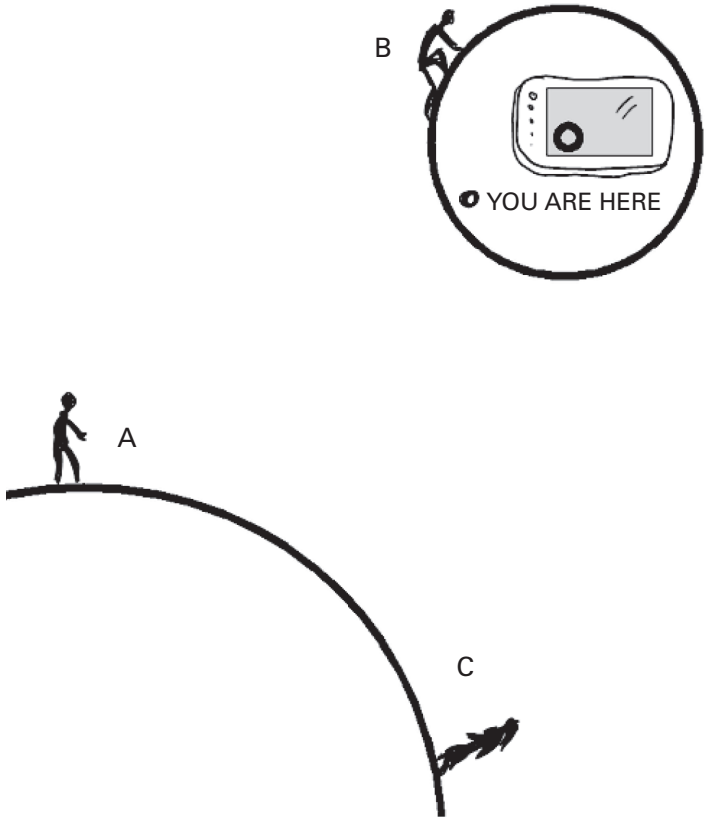

Figure 4.3
The position of A is tracked by a device whose screen is visible on the surface of the Moon. A is not within close sight range and can only use the display for navigation if she knows that the device tracks her position. B is within close sight range, but the device is not tracking his position. C neither is within close sight range, nor is her position tracked by the device. It is because of the possibility of many cases like B and C that we need epistemic segregation. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

within *close sight range for the map user*. Its character can be articulated as "whoever sees this from close up is close enough to the place represented by the Red Dot." It follows that the Green Dot of our example cannot legitimately bear the label "you are there." If it did, *Paula* would be misled into believing that the map represents *her* current location, and so would bystanders in Greece looking at the Moon. The correct label for the Green Dot, if Paula is to use it to help John in Rome, is "*John* is at the place represented by the Green Dot." In Paris, the map of Rome, although indexed or tracked, is not a situated map of Rome. It looks as if situatedness is the best and most cost-effective way to guarantee *knowledge* of situatedness; knowledge of situatedness creates *epistemic segregation*, a sort of protection from visibility from other, unwanted locations.

But epistemic segregation can be obtained by other means than through situation. We need to make sure that the map on the Moon was clearly labeled specifically for John and no one else. Using someone's full name and date of birth could sufficiently restrict the potential ambiguities in the label on the map (i.e., "This dot only tracks the position of John O'Toole, born in Berlin on May 8, 1969, and of no-one else"). John, and the other viewers of the map, must possess the semantic competence about John's name so that only John should feel entitled to use the map projected on the Moon for his whereabouts.

As epistemic segregation is what matters most, the fact that it can be achieved by means other than situatedness limits the theoretical weight put on situatedness. In chapter 1, we have seen that the mind is better conceived as not extended. Here, we make a claim to deflate embeddedness (situatedness). What matters is the exclusive, singular, and unambiguous usability of certain representations. Situatedness is a sufficient condition for exclusive, singular, and unambiguous usability, but not a necessary condition.

In other words: *Red Dotting = Green Dotting + Epistemic Segregation.*

However, not only must epistemic segregation be in place, but we must *know* that it is in place, that is, that we are the relevant addressees of segregated map information.

So maybe: *Red Dot = Green Dot + epistemic segregation + knowledge of epistemic segregation*

Perhaps surprisingly, situated maps (either static, i.e., nailed to the location they represent, or dynamic, i.e., GPS assisted and continuously updated as

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

the user moves) are conceptually less abstract than non-situated maps (e.g., the standard foldable paper map one may carry in one's pocket). The latter are best conceived as abstractions of the former. This means that, interestingly, the more abstract here preceded historically the less abstract. Note that conceptual priority is not automatically a form of priority in use.

Epistemic segregation is one of the poorly explored cognitive advantages of public representations. Epistemic segregation's dual is *orientability*. A fixed map is epistemically segregated, whereas a portable map is orientable. The non-situatedness of portable paper maps usefully contrasts with the *viewpoint relatedness of photographs*. Photographs inform us about both the scene they represent (e.g., trees or a rock) and the viewpoint the scene is represented from. Viewpoint relatedness is not situatedness: viewpoint relatedness is a *representational* feature of a photograph. It is made available from the content of the photograph, even though it does not represent the viewpoint.

(Distraction: Self-portraits in the mirror represent both the sitter and her viewpoint.)

Situatedness, as we have seen, is a nonrepresentational feature. Sure enough, some pictures can be situated on top of being viewpoint-related. For instance, the panoramic views at some observation points (figure 4.4) are situated in the viewpoint from where the represented scene is depicted, but this is not the norm for pictures. Most pictures and photographs end up in a place different from the one where they were taken.

As we saw in chapter 1, we can use photographs for navigation if we limit navigation to the area that is both included in the photograph and close enough to the viewpoint the photograph is taken from. It becomes evident that the cognitive advantage of maps over photographs is their *portability*, which is a consequence of their viewpoint-unrelatedness. Maps can be situated, and hence can be used for navigation, in many different points of the territory they map. Photographs used for navigation can only be situated close to their viewpoint in order to be useful. *Many inferences must be drawn to use a photograph for navigation in places that are not close to the photograph's viewpoint.* The inferential costs of using (non-satellite—see the next section) photographs for navigation grow as one moves away from the point from where the photograph was taken: the relationships between features displayed in the photograph must be computed and recomputed. The inferential costs of using maps for navigation are uniformly spread over the territory represented in the map. The costs of situating oneself on a map are high everywhere. The costs of situating one's point of view in a

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

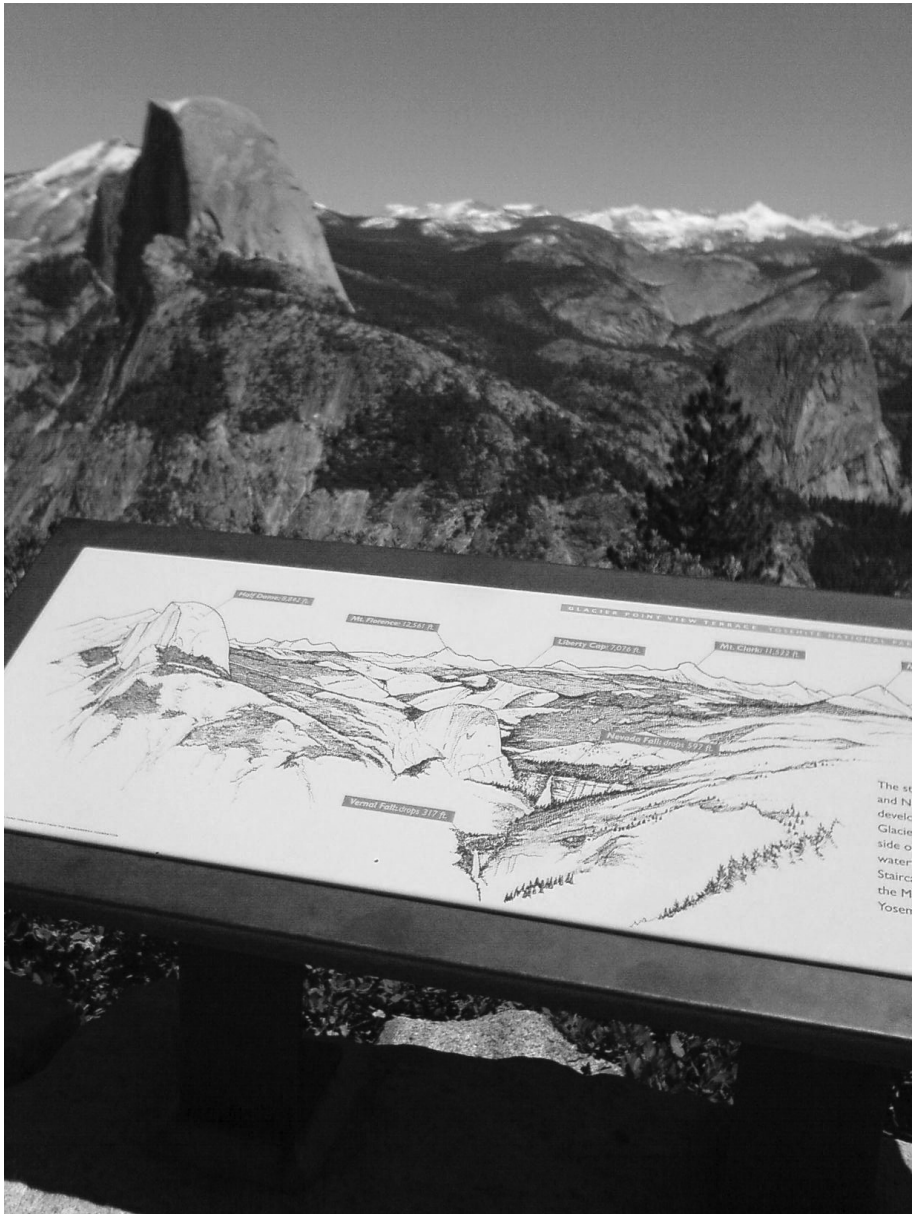

Figure 4.4
A situated picture. Orientation table, Glacier Point, Yosemite, April 2014. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

picture diminish as we come close to the point from where the picture was taken.

Toolbox: Point of View versus *Direction* of View

It can be argued that, strictly speaking, maps are *not* viewpoint unrelated. The case for non-unrelatedness can be made intuitive precisely by observing that maps can be superposed to satellite photographs. To address this issue, we can specify the distinction between viewpoint and direction of view in our treatment of axonometries (chapter 3). A viewpoint is situated somewhere along a *direction of view*. Maps are akin to axonometric (parallel) projections. They abstract from a viewpoint but preserve a direction of view. A map of Rome does not contain information about a specific viewpoint, but surely it excludes all viewpoints that are not along the direction of view (Casati and Giardino 2013).

(Distraction: We are considering here only maps of relatively small areas. Issues related to the curvature of the Earth and to the classical projection trade-offs of maps are marginal in this context.)

North-Up versus Target-Up versus "End-of-Next-Rectilinear-Stretch-Up"

Some like to orient maps north-up, others like maps to be target-up, and some even like maps oriented "end-of-next-rectilinear-stretch up." When it comes to orienting maps, cognitive advantages are weighted differently, or maybe map orientation is prey to old habits. A couple of things are of interest here. If drivers appear to prefer "end-of-next-rectilinear-stretch up" representations on GPS-based navigators, in the norm the preferred map display is picturified and perspectivalized.[3] In this display, it looks as if we are flying above our car, heading toward the end of the next rectilinear stretch of our path. Now consider old-fashioned, paper map hikers who do not enjoy the constant updating the GPS-based navigator can provide. In the "end-of-next-rectilinear-stretch-up" mode, the map is always oriented relative to the environment, and the user turns it around as the path unwinds. In the north-up mode, the user keeps her orientation relative to the map, but the map changes orientation relative to the territory as the path unwinds. In both cases, there are computations to be made. In the "end-of-next-rectilinear-stretch-up," we need to recompute our orientation

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

relative to the map in order to keep the map always oriented relative to the territory. In the north-up mode, we need to compute directions by looking at the map as it constantly changes orientation relative to the territory. Training can make a difference here.

Red Dot Maintenance

At this point, we are in the position to slightly hone our characterization of maps. Maps are things that on top of satisfying the (1) systematicity and dimensionality constraints, are such that (2) you can put a Red Dot over them, and (3) you can use them for navigating from your current Red Dot position to another position represented on the map.

In the following chapters, a number of representational artifacts will be seen as maps precisely because they incorporate these features. I will be relatively liberal with regard to definitions, as what matters here is our cognitive engagement with these items. Based on the features described previously, the underlying claim for each item addressed will be that from the fact one can use a certain artifact as a map, it therefore follows that the artifact is indeed a map.

Perspectival Anchors in Maps

Maritime maps give detailed information about the sea, particularly about soundings that are critical for landing or for avoiding beaching. The display of information gets denser near the coastline because the sea–land interface is critical. The display of territory inland is largely sacrificed and abstracted away. Exceptions are constituted by notable points such as landmarks that are salient for a navigator at sea, and particularly when approaching land. These are used, when visible, for determining the boat's position and calculating its course.

In a sense, all points registered on a map are perspectival anchors. In figure 4.5, we see those that are singled out by the map designer for being used as perspectival anchors.[4] In terms posed by Hutchins (1995), some "precomputations" are stored in the map; presumably one of these precomputations is a choice by the map designer to record salient points. In the light of our discussion on offloading in chapter 1, we can consider Hutchins's precomputation metaphor relatively harmless.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

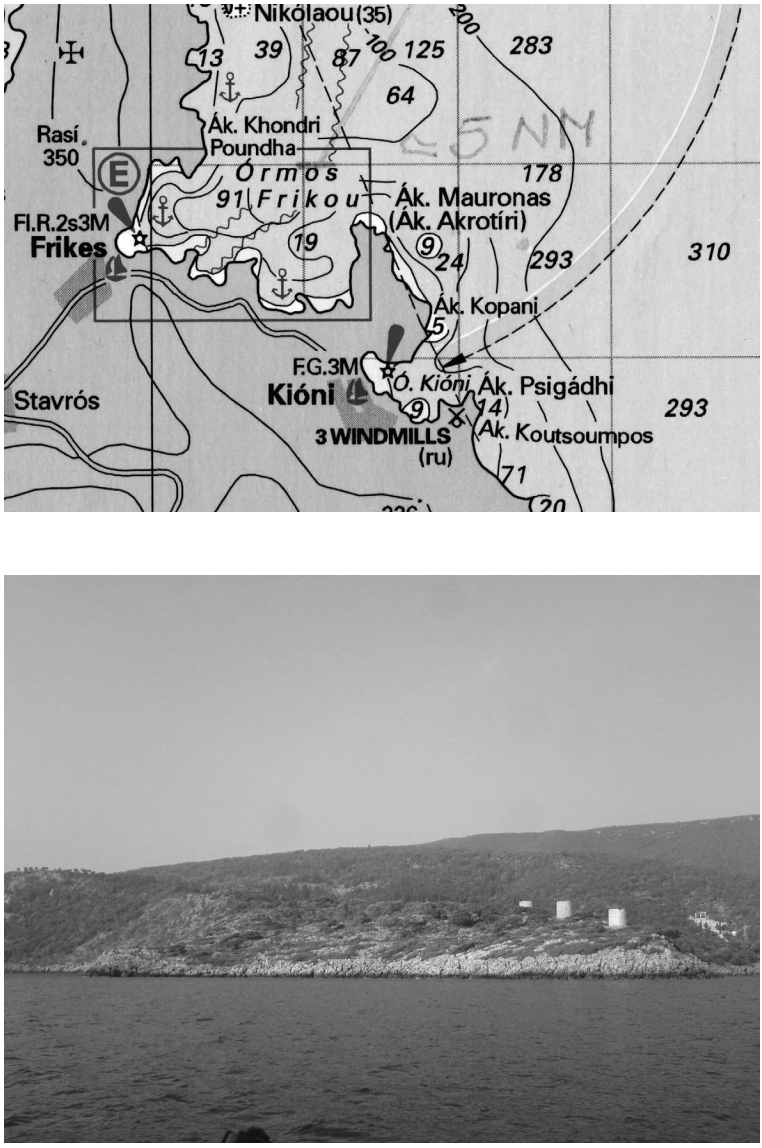

Figure 4.5
"3 windmills." (Top) Imray Chart of the Inland Sea, 2013 edition. (Bottom) A picture of the three windmills from the Sea (Image credit: RC, July 28, 2013). Bushes are visual noise to the sailor, except as their relative density enables her to detect a dent in the coastline.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Red Dots and Abstraction, Floors and Ceilings

Two public representations can be informationally equivalent but differ on extractability. Raising the ceiling (see chapter 1 and Kulvicki [2014]) is a way to make some pieces of information more extractable (i.e., freeing working memory that is no longer involved in extracting; no longer "decode first" so that resources are free for drawing other inferences). This works well for diagrams. We can graphically play on diagrams to make some higher-order properties of information salient. In the timetable example in chapter 1, we can color-code or space-code all names of trains between 9 a.m. and 11 a.m. so that the property of "being a train leaving between 9 a.m. and 11 a.m." is readily extracted. Extractability is a response-dependent feature. It depends on the consumer's needs and capabilities.

What about maps? There are many different consumer's perspectives on maps. One perspective is shared with consumers of diagrams. You may want to understand the relationships between certain land masses (i.e., which is larger? which is closer to the sea?). In this sense, one can raise the representational ceiling in many different ways (adding isolines, using the same colors for same properties, and so on). But another consumer's perspective is related to immediate action: moving from place a to place b. Consider two ways for making movement-relevant information salient. One is to put a Red Dot on a map. If maps are to work as navigational tools (i.e., for movement-relevant information to be easily extracted from them) that need to have nonrepresentational or epistemically segregated links to some of the places they represent, then raising the ceiling here is to have someone or some device set the Red Dot in the appropriate place.

Arguably, another independent way to raise the ceiling is to move from a bird's-eye view to a route perspective. You twist the map so you can foreshorten the path in front of you. But is this really a way of raising the ceiling? It depends on how you operationalize the notion. You definitely do not have to decode first. You free working memory, making it available for various inferences and decisions. But doesn't it sound odd to say that a map is made more abstract by being looked at the way a picture is? Or do we have to extend our notion of abstraction?

A proposal is that *uses* of different maps with the same degree of roughness can be more or less abstract. A map of the marsh area that is laid flat on the table is *used* in a perspectivally less abstract way than when the same

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

map hangs from the wall. When it is on the table, you can engage in navigational simulations that are closer to the physical situation (horizontal movements). When the map hangs on the wall, you do not take simulated movements on the map as instances of vertical climbing.

Maps and Territories: Self-Reference

Situatedness invites additional interesting questions. When maps are situated, are they on the map—that is, can maps be on themselves? Can a map be its own territory? Take for example figure 4.6 where the color gray = areas colored gray on this page (this is modeled after a meme about diagrams "gray = amount of gray on this chart"). Can a territory be its own map? For example, a line of people waiting to access a service could be taken as territory being its own map (chapter 9).

Every (portion of a) tiling of surfaces in the territory (even the tiling in your kitchen) trivially becomes a self-referential map if one so chooses. Tiling is possibly an example of a tautological map. But as trivial as it may appear, an orientation use can be found to self-referential maps (figure 4.6).

A different set of questions related to situatedness is created by Borgesian maps, which are neither self-referential maps nor self-mapping territories, but rather maps that are superposed to their territory. Self-referential maps and self-mapping territories can be considered as limit cases of Borgesian maps). Let's look at them in some detail.[5]

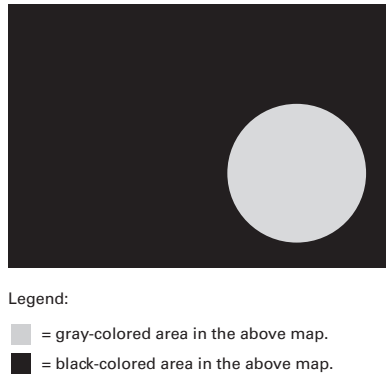

Figure 4.6
A self-referential map. It maps the distribution of color on the map itself. For doing so, it must put the color where it is. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Borgesian Maps

Borgesian maps are full-scale maps that are superposed on the areas they represent. They are often regarded as useless, absurd extensions of mapmaking practice.

> In that empire, the art of cartography attained such perfection that the map of a single province occupied the entirety of a city, and the map of the empire, the entirety of a province. In time, those unconscionable maps no longer satisfied, and the cartographers guilds struck a map of the empire whose size was that of the empire, and which coincided point for point with it. The following generations, who were not so fond of the study of cartography as their forebears had been, saw that that vast map was useless, and not without some pitilessness was it, that they delivered it up to the inclemencies of Sun and winters. In the deserts of the west, there are tattered ruins of that map, inhabited by animals and beggars, in all the land there is no other relic of the disciplines of geography.
>
> —Jorge Luis Borges, "On Exactitude in Science" (Borges 1944/1998, 325)

> "And then came the grandest idea of all! We actually made a map of the country, on the scale of a mile to the mile!"
>
> "Have you used it much?" I enquired.
>
> "It has never been spread out, yet," said Mein Herr: "the farmers objected: they said it would cover the whole country, and shut out the sunlight! So we now use the country itself, as its own map, and I assure you it does nearly as well."
>
> —Lewis Carroll, *Sylvie and Bruno Concluded* (Carroll 1894/1988, 169)

We argue that they are neither useless nor absurd and that they are quite common, following L. Frank Baum's vivid representation:

> "The road to the City of Emeralds is paved with yellow brick," said the Witch; "so you cannot miss it . . ."
>
> —L. Frank Baum, *The Wonderful Wizard of Oz* (Baum and Moser 1899/1986, 27).

They are, in fact, the conceptual grounds on which all other mapmaking is built. As such, they tell us some very interesting things about maps, the purposes they serve, and how they serve them.

A Borgesian map is a 1:1 map of the Empire that is exactly superposed on it. Borgesian maps were foreshadowed by Carrollian maps but differ from these in an important sense. Lewis Carroll imagined a world in which the cartographers' guilds were held in check, forced to keep their impractical maps folded up, so Carrollian maps are not superposed. Though Carroll

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

doesn't explain how they used their own country as a map of itself, perhaps the road to the City of Emeralds provides a clue.

Imperial ambitions are not prerequisite to Borgesian maps, though in those contexts things are apt to seem absurd. Being full scale and superposed are the only two features required, and since perfect superposition demands full scale and a specific manner of projection, our maps actually just require that. Humbler ambitions, like mapping one's desk, refrigerator, or circuit board, can thus lead to less ambitious Borgesian maps. In fact, Borgesian maps are legion.

Spray paint maps out a labyrinth of underground pipes and cables. A motherboard is marked to show its parts' connections. Airplane wings show where one may (not) walk. White blazes mark the Appalachian Trail from Maine to Georgia in the eastern United States. The town of Baarle-Nassau-Hertog, a patchwork of Belgium and the Netherlands, has borders marked all over the ground. Painted pavement that indicated the Checkpoint Charlie line between East and West Berlin was rescued from the bulldozers and rebuilt in the Mauer Museum. The line no longer serves its original purpose, but it was a prized artifact after the Wall's fall. One imagines some yellow bricks displayed in a museum long after the demise of the Emerald City, like bits of the Appian Way.

Not only are Borgesian maps far from absurd and useless, as actual examples show, but we claim that they are the conceptual grounds on which all other mapmaking is built. As such, they tell us some very interesting things about maps, the purposes they serve, and how they serve them.

The only two features that distinguish Borgesian maps from other maps are that they are scaled 1:1 and are superposed on the space that they represent. With that in mind, two points of clarification are in order.

First, Borgesian maps can be incomplete. A map scaled 1:1 and perfectly superposed on the area it represents can nevertheless be quite selective in which features it represents. If we were to insist otherwise, we would have to follow Josiah Royce (2012, 504–506) down a rabbit hole. He imagined a complete map of England, itself located in England. Such a map should include a representation of itself, which should include a representation of itself, and so on to Wonderland. Borges never suggests that the cartographers insisted on including all details on their maps. They sought coverage. Our account explains why seeking coverage, irrespective of completeness, constitutes an important ideal from which much mapmaking practice is derived.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Second, Borgesian maps need not be fully mimetic. All Borgesian maps are spatially mimetic. The representation of a road will have the same shape as the road because these maps are perfectly superposed on the regions they represent. But such maps need not be mimetic with respect to the properties they represent at these locations. A map can represent a road with a line of color, not a road. Maps can represent average rainfall across a region, but they do not represent average rainfall by instantiating average rainfalls. The distribution of pipes and cables beneath the street can be represented by patterns of spray paint. The yellow of the Yellow Brick Road can represent the way to the Emerald City, and be located along the way to the Emerald City, without *being* the way to the Emerald City. Carroll imagined people using their country as its own map. This seems profoundly unhelpful. Each feature and perhaps each thing is understood to represent itself, right where it is. This is not a representation that serves any obvious purpose. We take Borgesian maps—1:1 and superposed—to be useful and to constitute an interesting limit case of mapmaking practice. Since maps are parts of the world, they are, trivially, examples of people using parts of their landscapes as maps of those landscapes, but they are representations: selective, and not necessarily mimetic. In the summer of 2018, dry, hot weather in Great Britain caused crop marks, patterns of dry grass that correspond to the presence of ancient walls and mounds. Such natural occurrences are not, in and of themselves, maps. But those who understood the significance of the patterns could use their own landscape as a 1:1 superposed map of buried structures. An archaeologist could have created a Borgesian map indistinguishable from the crop marks, marking the surface selectively with weed killer. The marks would carry representational significance, just as the space in which they are arrayed would.

Useful Borgesian Maps

The world's cities contain Borgesian maps of laws that regulate space. Streets are marked to indicate permitted direction of travel, for example, for cars and bicycles, lanes limited to turning, or high-occupancy vehicles. Parking spaces, loading zones, standing zones, crosswalks, fire lanes, and so on occupy every available patch of road. A street's status is decided by the city council, and the decisions are marked in 1:1 maps with lines and signs superposed on the city. The indicators are not fire lanes and loading zones.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

They just stand for such properties and show where they are instantiated by being superposed on them. Cities not only regulate every patch of road but also typically mark every patch of road to indicate how it is regulated. Of course, streets are not exemplary media for maps. We usually think of maps as drawn on paper or computer screens. But, fundamentally, maps are representations that are, inter alia, marked arrays. The city streets fit the bill for being, inter alia, marked arrays.

Construction in developed spaces requires a careful touch, and so before work begins, roads and sidewalks are spray-painted to indicate the presence of electrical lines, steam pipes, and the like, lest bulldozers destroy them. Motorists wouldn't find these maps important, of course, because their actions are independent of the presence of these things. For construction crews, however, such maps are essential (see figure 4.7).

When the borders between countries are as complicated as the borders between parking spaces and loading zones, as in the towns of Baarle-Hertog

Figure 4.7
Street markings are maps superposed with their territory. Paris, 2017. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

(Belgium) and Baarle-Nassau (Netherlands), it pays to paint the borders onto the streets. Derby Line, Vermont (United States) borders the town of Stanstead, Quebec (Canada). Without a natural boundary like a river, the towns must mark the streets. The Haskell Library is in both towns, and a protected building in both countries. The floor is marked with the border (figure 4.8.) Again, the land around the United States/Canada or Belgian/Dutch border is not a standard medium for a map, but we do mark such

Figure 4.8
Haskell Free Library, Derby Line, VT, and Stanstead, QC. https://commons.wikimedia.org/wiki/File:HaskellLibraryBorderLine.jpg

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

surfaces with representations telling us about what that land is like. Though not exemplary maps, they are, we claim, examples of Borgesian maps that are very useful and far from absurd.

Large, complicated buildings like hospitals and parking garages often include Borgesian maps of themselves. A red line shows the way to "Emergency," a blue one to "Radiology," and so on. Green marks level 6 as a way of helping you find your car. Trails in the wilderness are maintained by repetitive use and occasional brush cutting, but they are also often marked with blazes. Because trails intersect, and it can be a disaster if you take the wrong one, we need markers to tell them apart.

One caveat, in some cases augmented reality (AR) can make use of virtual Borgesian maps, even though such virtual maps are just that: virtual. Point your camera at a spot, and you can see the kind of road it is, where it leads, something about its history, and so on. Another way to think of AR is that it reveals an otherwise invisible and ever-evolving Borgesian map, but it's not really there so much as virtually there. There are no marks in the space constituting a map. The British Geological Survey's Project Iceberg seeks to map the subterranean layers of London and might be implemented via an augmented reality map. Point your phone at the street and see indicators of what lies beneath. A similar point can be made about auditory variants on this, for use both by the blind and by the sighted. These devices announce locations as one comes to them, often as a way of giving directions. In many ways, such virtual maps give many practical advantages over true Borgesian maps, but they are derived from them. Whether one wants to call these AR artifacts maps or mere virtual maps depends on how fastidious one is about what constitutes a mark at a place.

Situatedness and the Conceptual Priority of Borgesian Maps

We have seen, discussing situatedness, that if a map represents a region that you happen to occupy, then, trivially, at least one point on that map represents the point you occupy. It does not represent that point as the point you occupy, however. To do this, a map must use Red Dots to indicate where one is and, sometimes, how one is oriented with respect to the surrounding space. These are situated, oriented maps.

Situation and orientation markers have an interesting property. The parts of a map that the markers occupy represent specific parts of space,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

and the Red Dot marker says of some represented part of space, that it is here. So, a map says, of a large region, that it is thus and so. The Red Dot, along with the map on which it is placed, says that here is thus and so. In that sense, the Red Dot bears an interesting nonrepresentational relation to the space: it is there. Of course, no matter where a marker is, it is there. So, for the Red Dot to work it must be coordinated with the part of the map that represents the space occupied. Similar thoughts apply to orientation.

One way to make this work is to fix a map's location, which allows one to place a marker on it that says what that place is like with respect to the others represented by the map. That is, one can locate a map in space, fix its location, and then the place marker on it, which shows where on the map *here* is. Examples of this are legions throughout the world's cities. Another way to make this work is to coordinate the GPS with a map that allows such markers to be placed on it dynamically. In both cases, one can add an orientation marker to indicate how a user is oriented with respect to the location she occupies.

Borgesian maps both represent a space and are perfectly superposed on it. As a result, every point on such a map both represents a location and occupies it. If you are at a specific location on a Borgesian map, you are eo ipso at the location it represents. Such maps need no situation markers since they come pre-situated. By virtue of this, all locations on such maps say, "here is thus and so." Perfect superposition also means that the map is oriented with respect to the user. I am here on the map and over there on the map corresponds to over there in the represented space. These maps are thus pre-oriented as well. Notice, by the way, that this feature of Borgesian maps is independent of how much the map says about the relevant place. One might imagine an ideal map to be complete, but the point for now is that the Borgesian varieties are special independently of whether they are complete. They can say very little, but whatever they say, they say of the places they are located.

Ordinary maps of many kinds are abstractions from Borgesian ideals. To see this, let's consider a simple example. You stand in a room whose floor is marked with a Borgesian map of the pipes, wires, and joists beneath it. Now imagine staying where you are but translating the map one meter in some direction. The new map is still useful, but you will need a dynamic "you are here" marker to use it. The marker will be one meter from the spot

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

you occupy. Move around, and the spot shadows you, one meter away. The spot shows where you are, given the way the map represents space. You are obviously not standing on the "you are here" spot, so in that sense you are not there. Alternatively, imagine starting with the Borgesian map and rotating it 20°. Again, the map is still useful, but you will need an orientation marker to use it. The marker will be displaced 20° from your orientation. Turn around, staying in the same spot, and the orientation marker moves with you, offset 20°.

Now imagine moving from the spot you occupy on the rotated map, and you have, in effect, both translated and rotated it. In this case, you need both the location and orientation markers to use the map effectively. In that case, there are few practical benefits to the map being at scale. Why not shrink it to fit in the palm of your hand? In this way, standard, situated, and oriented maps are abstractions from Borgesian maps. One final step. Remove the situation and orientation markers, along with scale, and you are left with an ordinary map. All maps are abstractions from Borgesian maps.

Only recently has the technology developed allowing for the dynamic situated and oriented maps we use on smartphones. Exemplars of maps from the past are typically those that are portable and thus without any situation and orientation markers. These maps are at some remove, logically speaking, from Borgesian maps, but their historical prominence, and perhaps priority, should not distract from the logical point made here.

Extensions of This Idea

It is far from obvious that historians thought of Borgesian maps as maps, and it is thus possible we might find more evidence of their use now that we know what to look for. The Romans used mile markers on their roads and occasionally other signs (Salway 2005). The markers indicated who funded the road and the distance from a given place. Sometimes they included the name of the place from which distance was measured, sometimes not (Baldwin Bowsky and Niniou-Kindeli 2006). Roman tabellaria (Salway 2001) were special markers that listed routes that radiated out from the places at which they were located. These tabellaria might have influenced the production of what we think of as ordinary maps (Salway 2005).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

In late-seventeenth-century Paris, under the Sun King Louis XIV, Nicholas de la Reynie ordered the first systematic attempt to light the streets of a city at night (Tucker 2018). Not every street in the City of Light had the lanterns, however, and the selection of such streets could also have served a mapping function. Illuminated streets are the main streets. We are not claiming that these are clear cases in which Borgesian maps were used. They are candidates for which some more historical investigation is needed. This case, too, is one in which people might use the presence of light as a map, even though we are reluctant to say that the light is a mark on a map. In that sense, this example is a bit like the crop marks mentioned earlier.

Though maps rarely show things in perspective, there are kinds of Borgesian maps that are closely related to projection. Mountaintops often feature maps of the peaks around them. They are shown in silhouette and labeled, often with indications of their distances and heights. Imagine, for example, the following picture set in place in front of the Churfirsten range near St. Gallen in Switzerland (figure 4.9).

This might seem more like a picture than a map, and it hardly seems to scale or superimposed. But remember that we are imagining this map to be fixed in place in front of the mountains that it represents. It gives a scale representation of the projective features of those mountains from the position a viewer occupies when viewing the map. This is a Borgesian map of the 2D layout of peaks from a point and is actually quite close to the augmented reality examples discussed previously. Point to a bit of the map, and, given that this is done well, you are pointing to the peak in the distance. The map is to scale, if the norm for scaling is a projection of

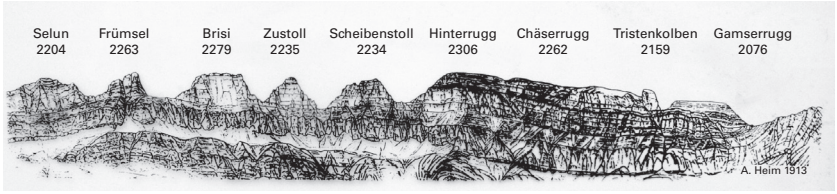

Figure 4.9
Darstellung der Churfirsten (Heim 1913). Public domain. Image: Adrian Michael. CC-ASA 3.0.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

the distant peak onto the screen of the map surface. If the foregoing seem like maps, then it might be that ordinary pictures can be Borgesian maps as well. Produced in real time, and oriented properly, pictures not only are maps of the Albertian Window, but they map the two-dimensional pattern from a point in a Borgesian way.

Maps are typically interpreted as representing how things are *now*, unless they are explicitly marked as representing some other time. Just as a Borgesian map is perfectly superposed over the space it represents, it is superposed in time. After the renovation, the extensive Borgesian map of pipes on the wall is inaccurate, unless we reinterpret it. There are two salient possibilities for doing so.

First, if we detach it from now and treat it as a map of how things are at time T, it can be accurate but not Borgesian, because it is detached from time. The new map says, "At time T, things are like this," but it doesn't say anything about when time T is (now, earlier, in the future, and so on). In that sense, it is detached from time. Second, we could reinterpret it as saying how things used to be. That is, it says now things are such that in the past they were like this. Because this map speaks to the current state of things, and is otherwise to scale and superposed, it is Borgesian.

In cities, parking regulations stay the same but their force changes over time. No parking here from 10:30–2:30 Tuesdays and Thursdays, but otherwise it's fine. The signs that map out regulations say this, but they do not change, as signs, over time, to mark the current force of the regulations. So, to the extent that they are Borgesian maps, they map the timeless regulations: either parking is permitted here or it is Tuesday or Thursday between 10:30 and 2:30, in which case parking is forbidden. We cannot generally control how we pass through time in the way we control passage through space, so it's unsurprising that time is not as relevant to maps, Borgesian or otherwise, as space is.

Conclusions

This chapter dealt with some of the *navigational* inferences we can draw from maps. Situating a map is a sufficient condition to use it for navigation, and we saw how epistemic segregation can be thought of as a necessary condition for navigational map inferences. We contrasted situated maps

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

and situated photographs and pictures in general, and then considered two limit cases of a situated map, self-referential maps, and Borgesian maps. We then contrasted Borgesian maps with Carrollian maps and ordinary maps, projections, and GPS-assisted navigation devices. The conceptual strength of a 1:1 superposed representation is shown by its ability to generate all cases of maps by abstracting from some of the Borgesian map features. In itself, the category of Borgesian maps remains an important descriptive tool that unifies a number of representational genera.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

5 Maps and Models

Models, Globes, and Maps

In previous chapters, we hinted that there is a porosity in the distinction between maps and models (chapter 2). Here we further discuss analogies and differences, focusing on globes that represent the Earth. Globes are 3D *maps* of the Earth, used for various purposes, nowadays mostly in instructional settings. Built in various scales, globes are typically manufactured with a support that carries information about the Earth's situation relative to the Sun (i.e., the axis' inclination relative to the plane of the Ecliptic).

Calling a globe a map as opposed to a model requires some preliminary arguing or terminological clarification. Earlier, we discussed the fact that there can be 3D maps of 3D territories. Ordinary use reserves the term "map" for 2D representational vehicles and the term "model" for 3D representational vehicles. However, some 3D representational vehicles are not thereby models (artistic sculptures, for instance), and some 2D representational vehicles are models (such as a tailor's cutout for a dress shirt). It is safe to assume that what turns a map into a model is not their respective syntactic features but rather a certain *use* one can make of the map. If we limit ourselves to concrete models and exclude abstract mathematical or computational models (Weisberg 2012), then a model is a certain physical item that is interpreted as representing a certain object or situation in the world and is designed and used to make observations, and to generate or draw *inferences* about the represented situations in the world. This prima facie characterization makes room for maps to be used as models. By looking at a map and interpreting it, you can make inferences about the reality represented by the map. And, as we have seen in chapter 2, maps can definitely be created with the specific intention to be used to draw inferences from and about a map's contents.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Classifying Models

I provisionally accept the basics of Weisberg's (2012) classification for expository purposes, distinguishing between concrete models (a 3D model of a dam), computational models (game of life simulations), and mathematical models (equations representing a collision). He specifically targets scientific or engineering models. Other functions and uses of models that are absent from this range include pedagogical models (a globe), appearance models (typically used in architecture: what will your house look like?), reverse engineering models, and tailors' models. This neglect of other types of models may have an impact on the large project of a classification of models. For the sake of the present discussion, these complications will be omitted; but we'll have something to say about some of these items.[1]

Earth Globes: Naked

A globe is a 3D map (or a 2D map embedded in 3D space—the difference being of no importance in the present context). I will talk of *naked globes* when I discuss globes without supports and without other representations of Earth's extrinsic features, such as Earth's axis in relation to the Ecliptic. A naked globe represents various physical features of the Earth's surface without representing any relationship between the Earth and items that are situated outside its surface. A study of a naked globe allows us to draw countless inferences. Some of these inferences are related to the map-like structure of the globe (see chapter 2). By looking at a globe, you can discover that Madagascar is an island (a 2D feature) and that Italy and New Zealand are antipodal (a feature related to the 3D embedding). You can predict that you can go from Lisbon to Athens by land, but not from Lisbon to Antananarivo. You can also predict that it will presumably[2] take you longer to fly from Lisbon to Tokyo than it takes you to fly from Lisbon to Paris on each respective shortest route. These facts and inferences are readable from the syntactic features of the globe: the shape of certain regions, the parthood, topological and metric relations holding among them, and their 3D organization.[3]

In the same sense, a naked sculpture of a bear (figure 5.1) would represent just the bear and the relationships between parts of a bear, without specifying their relations to any other items that are external to the bear itself. As such, a naked sculpture of a bear can be situated in any orientation

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 5.1
A sculpture of a bear, decontextualized. Unidentified author, Inuit art, 1990s. Author's collection. Image credit: RC.

relative to the ground. By looking at it, we would not know if it is the right-side-up sculpture of an upside-down bear, for instance, or the right-side-up sculpture of a right-side-up bear (or any other combination for that matter). It may even be the sculpture of a flying bear. Of course, bears do have a *right* right-side up (or spend enough time in a position such that ultimately we consider that position as a normative *right* right side up for them). On the other hand, there is no normative right side up for the naked globe, except in a weak, irrelevant sense dictated by tradition—most of the northern hemisphere on most ordinary globes faces away from the ground when the globe is placed on its axle.[4]

Integrated Globes

Integrated globes build upon naked globes and add to naked globes some syntactic features that relate the representation of the Earth to representations of some celestial features. An ordinary globe that is purchased in stationery shops is typically mounted on a support that is used to indicate the spin of the Earth's axis and the tilt of the Earth's axis with regard to the Ecliptic. The representational vehicle integrates the globe in a larger representational space. It provides the globe with an axis, a now salient syntactic feature, and relates the axis to the Earth's plane of revolution. Integrated

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

globes differ from naked globes in terms of their syntax, and hence in terms of the additional inferences they support. In this case, integrated globes unlock inferences about some of Earth's extrinsic features. If by studying a naked globe you can see or infer that Nebraska is a proper part of the United States, or that Madagascar is an island, or make inferences related to the 3D embedding of the map (such as that New Zealand and Italy are antipodal), by studying an integrated globe you can understand more complex, exciting facts about the Earth. For instance, the globe's spin simulates the Earth's spin. (An escapement mechanism could constrain the spin in order to maintain the same direction as that of the Earth's rotation on its axis.) When the pedestal of an integrated globe lies on a horizontal plane and the globe is related to a situated light source representing the Sun, the integrated globe can be used for inferring various facts about the location of the Earth and the Sun during the year. The integrated globe is able to show these changes relative to the orientation of Earth's axis because it keeps its tilt during a revolution of the Earth—shadows on the globe become another syntactic element of the vehicle. The addition of other syntactic elements, such as a representation of the Moon, provides a representation of the mechanism of both solar and lunar eclipses.

Situated Globes

I would like to add to our list yet another category of globe models, that I'll call *situated globes*. I'll spend some time discussing the category as it provides a fascinating illustration of a number of principles governing map use in situation. Situated globes are real-time, situated, analog representations of some features of the Sun-Earth system. In the simulation they enact, the following holds:

- The globe's axis is parallel to the Earth's axis.
- The region of the globe representing the place the globe is at always sits on top of the globe.
- The *actual* Sun, not a representation thereof, is one of the syntactical objects along with the light the sun emits.
- The passage of time represents the passage of time at the 1:1 scale (one day equals one day).
- *Now* in the model represents *now* in the world.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

These features articulate the globe's being *situated*. As a consequence of these features' holding, the globe's spin on its axis has the same period as the Earth's spin.

To set up the situated globe and have it work, one has first to deconstruct it, separating it from its axle. Second, one should find the representation of one's location (of the globe's location) on the globe and have the line connecting the globe's center to this position point to the zenith (figure 5.2). Third step, one must orient the globe so that the direction of the meridian one is on is parallel to the direction of the corresponding meridian on the globe, with north in the same direction in both cases. Finally, one must simply nail the globe to its position and leave it alone, so that it cannot move anymore relative to the ground.[5]

To an observer on Earth, the globe may appear static as it does not move relative to the ground. However, the Earth spins, *and hence the situated globe is actually in motion in its larger astronomic environment*. Carried along with the movement of the Earth, the situated globe changes its situation during the day relative to the Earth's and the Sun's direction. Relative to that direction, the situated globe spins on its axis and turns on itself in (about) twenty-four hours. Given its orientation, the pattern of light and shadow on its surface corresponds closely to the pattern of light and shadow on the surface of the Earth at any given moment. This in turn means that the model only "works," for a given location, during daytime at that location. At night, the globe is in the shade of the Earth and cannot mimic the light distribution necessary for globe-to-Earth comparison. In a sense, the situated globe is a kind of sundial. Equatorial sundials, as we discuss later in this chapter, are straightforward abstractions of the situated globe.

Observations made on the situated globe allow the user to get facts about the current situation of the Earth (figure 5.3). By looking *now* (in daytime) at the situated globe, one can learn where it is night *now* on Earth and where it is day. Whenever the situated globe's terminator (i.e., the line separating light from shadow) meets the representation of a city on the globe, the city in question meets the terminator on Earth. If shadow follows light, then the Sun will be setting on that city at that very moment. If light follows shadow, then the Sun will be rising there. You can find the spot on Earth where the Sun is at the zenith (i.e., in the center of the lit-up area). The shadow of a stick that is normal to the surface

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 5.2
A naked globe (left) represents the Earth but does not have syntactic marks to indicate its relation to the Ecliptic. An integrated globe (right) adds those marks: a pedestal and a fixed tilt of the axis. To turn an ordinary globe into a situated globe (bottom), find your current location on the globe (here: Hanover, New Hampshire, United States) and make it the top relative to your environment; then align your meridian and its representation on the globe so that the nearest pole is in the appropriate direction. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 5.3
The aspect of the situated globe of figure 5.2 evolves as the day unfolds. Observe the concomitant evolution of the local pattern of light in the garden. Etna, New Hampshire, Summer Solstice 2014. Image credit: RC.

on either pole on the globe indicates the meridian where it is currently local midnight on Earth. (Local noon and midnight are, in general, different from normalized noon and midnight.) The continuation of that meridian opposite to the shadow's direction indicates places where it is currently local noon on Earth. Each pole will be shaded from the Sun for half of the year, while the other pole will be at the same time lit up by the Sun. This dynamic remains the case except at equinoxes, when both poles are in view of the Sun. For most locations where the globe is situated, its axis—which is parallel to the Earth's axis—will be tilted relative to the surface of the Earth. This means that the hemisphere opposite to the hemisphere containing the representation of the place where the globe is situated will be lit up during part of the year (the local winter) in a possibly surprising way (i.e., it will look as if light comes "from below" onto that hemisphere).

Cognitive Advantages of Situated Globes

The cognitive advantage of the situated globe over the integrated globe is that the simulation's kinematics is not only completely controlled by the physical situation it both represents and is embedded in. This means the simulation uses the Sun instead of a representation of the Sun and uses the Earth's rotation instead of a simulated movement to ensure a reliable representation of Earth's shadow pattern. The shadow pattern on the situated globe is however interesting only insofar as it is a *representation* of the real earthly shadow pattern—and thus only as it can be used as a model of the latter.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Situated Globes, Sundials, and Abstractions

Most sundials can be seen as *abstractions* based on the situated globe. Their difference from the situated globe is that they may abstract from certain aspects of the 3D representation of the Earth that the globes provide. An equatorial sundial (figure 5.4) follows the same orientation constraints of the situated globe but collapses the 3D structure of the globe onto the equatorial plane.

Other types of sundials further abstract from the requirement that the shadow be cast onto a syntactic element that represents an Earth feature (e.g., the globe's surface or the equatorial plane). An extremely abstract and conceptually thin sundial is only sensitive to the relationship between the local situation to the Earth-Sun direction. This further abstraction turns sundials into time measurement devices that only track time.[6]

Limits to the Use of Situated Globes

Some limitations of the situated globe model ought to be taken into account.

1. The situated globe's axis would spatially coincide with the Earth's axis only if the globe was situated at either pole. This means that for all other

Figure 5.4
Equatorial sundials are situated globes collapsed to their equatorial plane. Rio de Janeiro, Botanical Garden, May 2013. Image Credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

locations the globe's axis at that location will describe a cylinder around the Earth's axis during the day. However, for most practical purposes, the distance between the surface of the generated cylinder and the Earth's axis is negligible compared to the distance of the Earth to the Sun. And given that the globe is rigidly tied to a rotating Earth, it is forced to spin at the Earth's spin period in any case.

2. The situated globe's terminator is shorter than it would be on a scaled model. A globe one meter across is a representation of the Earth at, say, approximately 1:12,500,000. The globe's terminator would represent accurately the Earth's terminator only if its distance from the Sun and the Sun's size were rescaled up the same factor. If the globe is intended to represent the shadowed portion of the Earth, then the shadowed portion of the situated globe is less extended than it ought to be. In this case, as in the preceding one, the effect is minimal.
3. As mentioned, the situated globe only works during the daytime. (Some borderline cases are no exception to the daytime rule of work, such as when a globe is located on either pole in the respective spring and summer seasons, where daytime lasts for half of the year.) This means that one cannot time the length of day and night with a globe at different locations on Earth.
4. Given the real-time anchoring, counterfactual simulations, or modifications of the temporal factors, are not possible to "run" on the situated globe unless one modifies some of its physical properties. It is an observational model.

What type of observations and inferences are made possible by the use of the situated globe? One can predict sunrises and sunsets. One can check the occurrence of a sunrise or a sunset at the location of the globe. One can find the current noon meridian. One can observe seasonal light and shadow patterns as they are at the time of observation. More complex operations can also be performed. By turning the globe upside down and reorienting it accordingly, one can have an image of what the seasons will be in six months' time, with time flowing backwards, from spring to winter (more about this later). In all these cases, one is displacing the computational burden onto the visual system, moving from linguistic inferences to visual inferences.

By construction, when seen from outer space, all situated globes look alike both in aspect and orientation relative to the Earth-Sun direction because every globe rotates in sync on an axis that is parallel to the Earth's

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

axis. This means that whenever the globe is sunlit, pictures of it taken at the same moment in, say, Delhi or Capetown will be indistinguishable from each other when it comes to representing the current pattern of light and shadow on Earth. The situated globe is a truly globalized model of the Earth. Its being situated makes it location-independent. If you are in Cairo, calling your scientific correspondent in Delhi to read out a present measurement on the globe at its current location will be trivial and useless: the measurements will be the same on your globe if both globes are presently sunlit. The difference will only be in each globe's orientations relative to the Earth's surface and, of course, in the local time at measurement.

Situated Globes and Faceless Globes

As a matter of fact, each sunlit spherical object on the Earth's surface contains the basic syntactic elements of a situated globe, barring geographic markings. Sunlit spheres such as balloons, basketballs, and the other like objects are situated globes that abstract from geographic representations of intrinsic features of the Earth. Each sunlit sphere can be looked at as a model of Earth as seen from outer space, as it is in relation to the Sun at the time of observation. (The Moon itself can be used as one such model. Looking at the Moon, one sees how the illumination pattern of the Earth looks from a position in outer space. Galileo already noted how the Moon and the Earth have a dual pattern of phases when seen from one of the two bodies.) Of course, when looked at under this abstraction, some regions of a given sphere are implicitly assigned an earthly referent that can be uncovered inferentially. One could use, for example, a soccer ball to say, "Given that it is noon here, and that that region faces the Sun, that region represents our location, and that other region represents the North Pole."[7]

Anchoring Situated Globes—Red Dotting Implicit in the Topmost Position

In chapters 2 and 4, we introduced some key features of maps and map use, and we stressed the importance of situatedness and epistemic segregation. The distinction is of some importance as we proceed to look at the functioning of situated models.

It is not essential that a physical model of tides in San Francisco Bay be located in San Francisco or in Tokyo, as the differences in gravity in the two

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

cities are marginal relative to the amplitude of the tidal effects one plans to model. But it does matter if a situated globe in San Francisco is oriented with the syntactic representation of San Francisco at the top, or at some other location instead. If the representation of San Francisco is not at the top, then the globe in San Francisco would not be oriented. The situated globe is anchored to the ground as is the static "you are here" map. The topmost syntactic element of the situated globe is functionally equivalent to a Red Dot. Its character would thus be "Whoever sees this from close up is in the vicinity of the place represented by the topmost part of the globe."

The situated globe can be used for navigation, but this is not the main point—sure enough, if the globe is not oriented it would not be possible to use it to navigate. I seek to underscore that if a situated globe is not oriented, then it does not *represent correctly* the various extrinsic features (astronomic relations) of the Earth that it is supposed to represent. The pattern of light and shadow on its surface would not correspond to any pattern of light and shadow on the Earth's surface. Orientation of the globe is thus a necessary condition for its *representational* capabilities. The correct setting of its orientation is in turn a necessary condition for its usability as a model (i.e., for drawing accurate inferences). Situatedness is a key feature of the situated globe. Short of situatedness, the globe cannot be used because it would not display correctly the properties that its real-time ambition dictates. Accordingly, epistemic segregation comes for free in this case.

The situated globe bears two independent situated components that are captured by the "you are here" character and by the "orientation" character, respectively. Minding epistemic segregation, the orientation character of the syntactic element constituted by the topmost surface of the globe can be expanded as follows: "*Whoever sees this from close up has a vertical direction that corresponds to the direction going from the globe's center to the location of its topmost part.*" If the globe is misoriented, the "you are here" character and the "orientation" character will inevitably diverge. You would not be located where the topmost part of the globe puts you, or you would not be oriented in the direction going from the globe's center to the "you are here" Red Dot. There is a trade-off between correct location and wrong vertical, and correct vertical and wrong location.

There are some interesting uses of a *misoriented* globe. In a limit case, you can do the following: find the correct local orientation for a globe, turn the globe upside down (Casati 2013), and sync the respective directions toward the North Pole at the surface of the Earth. In this case, the globe's axis is

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

parallel to the Earth's axis, but its north and south are inverted. The "you are here" syntactic element is now situated at the globe's area that is closest to the Earth's surface. In this configuration, the globe would represent the pattern of light and shadow that will occur in exactly six months, with time flowing backwards. (As the axis is inverted, the globe rotates clockwise if seen from above *its* North Pole.) Representations that use indexical elements can misfire in many ways. We have seen how John could not use the Green Dot map for navigation in Rome when the map was situated in Paris. The inverted situated globe cannot be used to directly ascertain what parts of Earth receive shadow or light, now or in a few minutes. Sunsets and sunrises are inverted on it.

The possibility of misrepresentation guarantees that the notion of situated models is not trivial. By the simple fact that they are concrete, most models are automatically situated. Each concrete model is by its very nature situated in that it is located somewhere in space-time. But the normative elements of situatedness depend on the representational scope of the model, which in turn depends on the intentions of the model maker.

The incorporation in the model of self-tokening items, such as the Sun in the case of the inverted globe or gravity in the case of a model of tides or waves, does not automatically generate situatedness for the representational element. Waves and tides can be generated wherever on Earth the wind blows and gravity exerts its pull. A model of the shadow pattern, on the other hand, requires fine-tuning of the location and orientation parameters.

The category of situated models allows us to hone our classification of items we use as models. I suggested that the situatedness of the underlying representation carries most if not all of the theoretical weight of the notion. The situatedness of the representational component is an important feature of many epistemic artifacts.

Appearance Models

We build physical models for many reasons (Weisberg 2012). We want to test the physical properties of a building or of a vessel; we want to optimize the form of the wing of an airplane. In all these cases, we do not care about perspective. What does matter is the intrinsic properties of the object that is modeled. However, a reason to build physical models is to show *how a certain object will look* once built on a larger scale, or from perspectives we do

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

not typically have. An architectural model of a building is typically created with this purpose (figure 5.5). We can play interesting make-believe games with appearance models. We can grasp not only what the building looks like from a certain viewpoint but also what it is to be in a certain position relative to the building. Through appearance models we can also see what it is to look at a model from a position in the environment around the model, what paths are afforded, and what a change in viewpoint can imply.

Appearance Models and Maps

In the context of explanation, communication, and education, much too simple objects are not sufficiently challenging, and much too complex mathematical formalisms prove discouraging. Appearance, perspectival, and situated models fulfill the need for "intermediate objects" in educational contexts. Figures 5.6 and 5.7 provide examples of situated models as both concrete and appearance models.

Figure 5.5
Navigation and change of perspective in an appearance model. Vicenza, Palladio Museum, November 2014. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

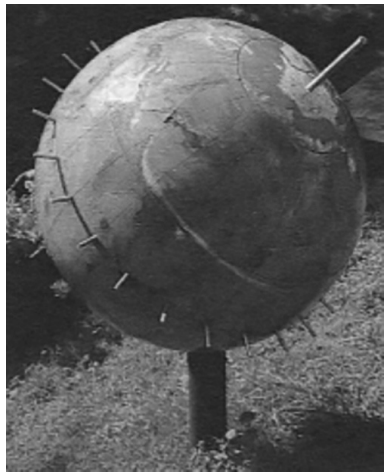

Figure 5.6
Moroder's sundial (Bolzano, Italy). Some syntactic features (posts on the equator) are added to the oriented globe to turn it into a sundial. It is possible to determine local noon on any meridian. The abstract ceiling has been raised. Image credits: Common domain.

Figure 5.7
What the Sun sees—an Earth without attached shadows. When on the picture the shadow of the situated globe is minimized, the shadow of the camera falls on the region that represents the current location of the globe (a place near Fontainebleau, France). Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

The Astrolabe as an Appearance-Computational Model

Astrolabes (figure 5.8) are perspectival devices. They offer their users a representation of the sky from a given place at a certain moment. Astrolabes typically are set for a certain latitude. By adjusting the *rete* (the mobile web carrying representations of stars) to match the time of day and the day of the year, you obtain a map of the sky as it is visible from that latitude at the specified time and day. You can match this map with the visual scene. This perspectivality of the astrolabe depends on the fact that celestial maps—not only the ones used on the astrolabe—in general are full-fledged maps of the visual distribution of stars. They are not maps of stars, as we saw in chapter 3. Celestial maps are conceivably the fittest candidates for a story about the origin of maps. In contrast with the impossibility of a bird's-eye view of a territory before human flight, the stars have always been in full sight and a celestial mapmaker could just inspect the celestial territory visually.

In the language of Kulvicki, the simplification/constraining of viewpoint is a way to raise the determinate floor, by making only a certain view available. This process in turn is a form of perspectival abstraction, as discussed in chapter 1.

Spatial Models versus Maps

Spatial models provide an interesting contrast class for maps. Spatial models, such as a model of the spread of epidemics (figure 5.9), do refer to specific geographic territories, are 2D, and are possibly analog. But spatial models are not maps as they do not satisfy the mereotopological constraint. For example, immunity spots are not to be found in the territory where you would expect them if the spatial model were a real map.[8]

A number of spatial inferences can be drawn from the spatial model of epidemic spread, but these are generic and do not involve any specific location. Spatial models bridge the map-diagram divide from the map side: they are born of maps, stripped of the locational precision of the properties they map. The river diagram of chapter 3 bridges the divide from the opposite side, as it enhances the spatial structure of a diagram.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

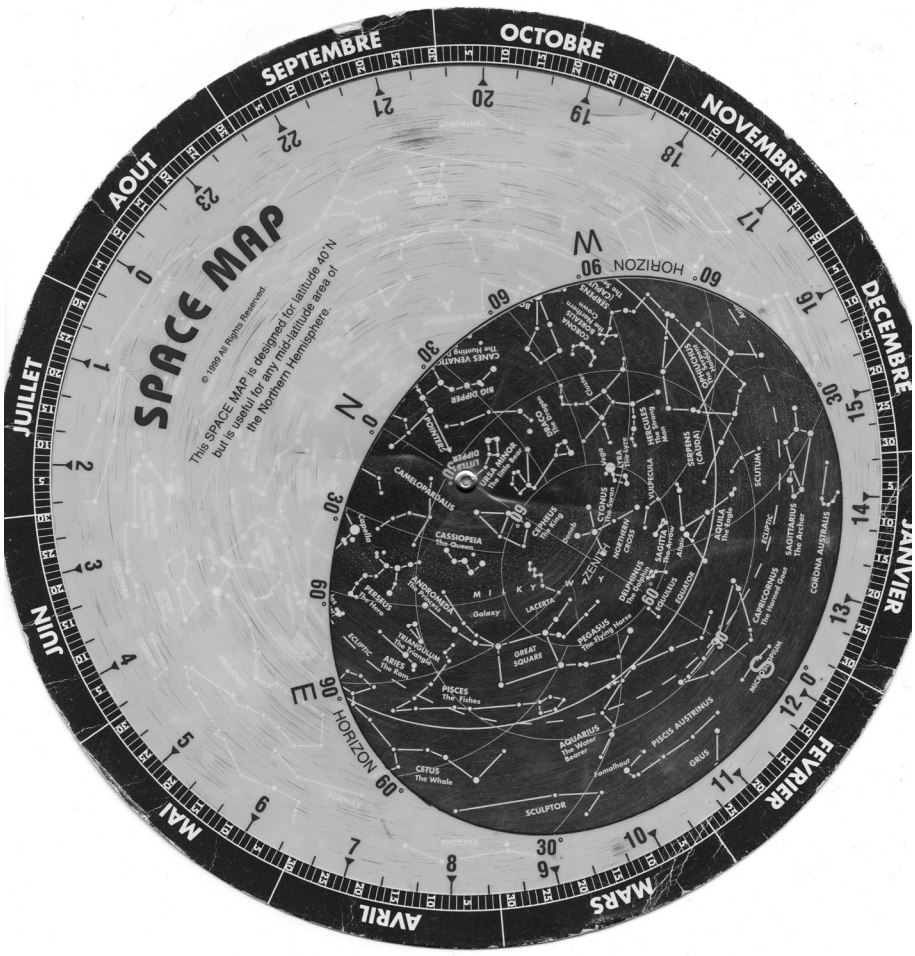

Figure 5.8
A modern astrolabe. By inputting date and time, you select from the star representation for a given latitude those stars that are visible at that given moment. In this particular display, this is the sky that you would see on July 8th at 2 a.m. from a latitude of 40° N. But it is the same sky that is visible on July 23rd at 1 a.m., and on August 7th at about midnight, and so on. Author's collection. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

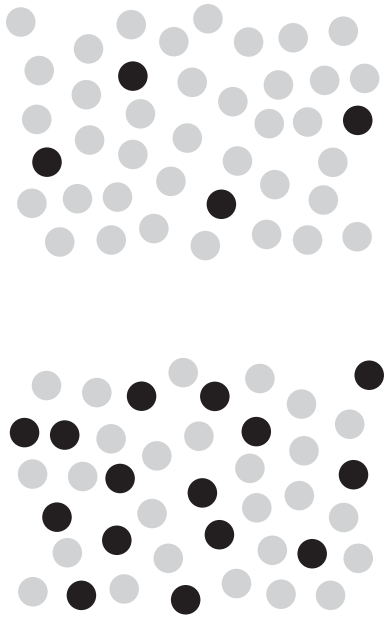

Figure 5.9
A model of disease spread can be both spatial and refer to a specific territory, without being a map of that territory. Image credit: RC.

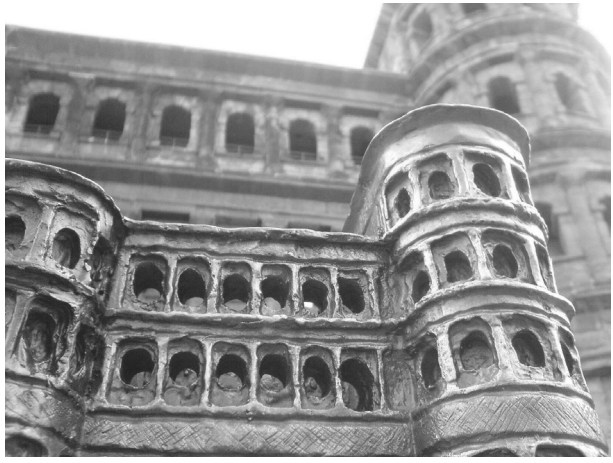

Figure 5.10
An appearance model can be situated. Trier, June 2015. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Conclusion

In this chapter, I made a number of claims. Models are individuated cognitively by their observational and inferential capability. This is a central design feature of models. Under the control and the intentions of their makers, naked, integrated, and situated models are distinguished by some of their syntactic features, which may integrate self-referential tokens, and by their embedding in an oriented space. Situated globes and sundials are real-time, situated, analog representations of some features of the Sun–Earth system. There is more to model taxonomy than simply representation, and situatedness is an important dimension for model classification. Appearance models need not be situated (although they may, as in figure 5.10), but they use the situation of the observer as an ingredient of the simulation they provide.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

6 Clockfaces

This chapter extends the map account to timekeeping devices.[1] Clocks display—working in Mode 4, a proxy mode—the current time on a (portion of) a temporal map that is a more or less finely tuned calendar, which itself (the calendar) is a map. In the past, we did Red Dotting on vintage, tear-off calendars, working in Mode 3, by tearing off pages day after day to keep track of elapsed time by situating ourselves on the calendar. Clock mechanisms enable Red Dotting on the clockface.

Timelines and calendars are abstract temporal maps. Their territory is a sequence of temporal events that unfold in time. Events referred to here include sunrises, sunsets, meridian passages, and the like. Calendars typically do not represent events in time, whereas timelines do. Calendars limit themselves to providing the empty grid on which representations of events can be placed. They are paced, a bit like the coordinate grid of an empty geographical map.

A clock is a device that places and maintains a temporal Red Dot on a calendar. The calendar may not be clearly visible in the clock. But some part of it must be accessible. A clockface, therefore, can be seen as the accessible part of the calendar the clock is navigating.

What Is a Clock?

A clock is an instrument for tracking and telling the time, and that can thereby be used as an instrument for generating true beliefs about time. Moreover, most clocks track time by first measuring and then counting time intervals, and finally displaying the result as an instance of a Red Dot, changing the place of a peculiar syntactic item, the clock hand, on a

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

calendar. In the class of epistemic instruments, clocks fall in the subclass of measuring instruments, together with rods and counters. And, further, clocks have this peculiar feature that they are tracking instruments.

Clocks have at least to say or indicate what the current time is, in a broad sense of the terms "say" or "indicate" (e.g., by the position of the hands on their face, by numbers on an analog display, or through a programmed bell-ringing device or loudspeaker that voices each hour change). Prima facie, a clock is something that says what the current time is, that is, it says what time it is at the moment at which it says that time. This is a token-reflexive indication: what is indicated essentially contains what performs the indication. In the vocabulary of the previous chapters, the clock is offering a Red Dot on a time map.

Time is measured starting from some given zero time, t_0. What a clock indicates is how much time has elapsed from t_0. The indication makes direct and token-reflexive reference to the present time and makes an oblique reference to t_0. The time map incorporated in a clock is thus like a long line whose points refer to temporal moments, and a clock aims to indicate on its face where the current time is.

----------x--------

Things that measure past (or future) intervals are not clocks. For instance, one could imagine a device that says right now how much time has elapsed between Napoleon's death and Bismarck's death. Such a device need not be a clock: for this device, you simply input a couple of dates and it computes the temporal interval between them. If you want this device to be a clock, then it has to indicate the time of the events it deals with when these events occur. It has to incorporate some token-reflexive pointer. Accordingly, a computer is not yet a clock (even though it may contain one); clocks cannot be simulated by Turing machines.[2] The same holds for measuring rods: they have to be used indexically. A machine that computes the distance between two points of different latitude and longitude, based on the input values of these coordinates, is not a measuring rod.[3]

Timers and Time Followers

I shall here distinguish between *timers*, *time followers*, and *clocks*. A typical example of a timer is the oven timer. You get it started by assigning it

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

a certain temporal interval. Once this interval has elapsed, the timer will stop (and usually emit a signal). A timer tracks the time, but it does not tell the time. You would not look at the oven timer to know what time it is. Certainly, one could use the oven timer as a clock (just as you can use your watch as a timer); there is, however, an important difference between the two objects. An analogy between the *spatial web* and the *measuring rod* will help to describe these differences.

Imagine that the Earth's surface has been covered with a web of lines, emanating from the center of Navona Square in Rome (which we will call $place_0$ or p_0) and calibrated in units of 1 km. Then to know what your distance is from p_0, you only need to look at the marker nearest to your present position and read off the distance.[4] You may also want to know what the distance is between your position and that of your neighbor, who is a bit closer to Rome on the same web line as yours. You can look at the line, but you can only read off your and then your neighbor's respective distances from Rome. To get the required information, you will have to *compute* the difference between the two. If instead you had a measuring rod at your disposal, you could have laid one if its ends at the edge of your house and read on the rod the number closest to your neighbor's house. The rod directly tells you the distance you are after (figure 6.1 and 6.2).

I am suggesting that *the rod is to the measuring web what the timer is to the clock*. A clock measures the distance from a fixed time t_0, while a timer measures any time interval.

(Distraction: The difference between rods and timers, i.e., that one can freely move the former in space whereas one cannot freely move the latter in time, need not bother us here.)

If you want to use a clock as a timer, you have to carry out a computation, whereas the timer carries out the computation for you. A borderline case is constituted by the use of ordinary watches for telling the hour of the day: these are clocks used as timers relative to midnights.[5] But if you want to use a timer as a clock, then you have to carry out an extra computation, where the clock normally spares you the computation. Thus, one can conceive of a clock as a timer that has been tuned or set, just as one could conceive of the measuring web as a set of fixed, tuned rods. (The timer spares you the computation, but is the computation offloaded to the timer? No. The timer does not compute anything. Once more, Turing machines cannot simulate clocks.)

One should make a finer distinction between stopwatches and timers.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

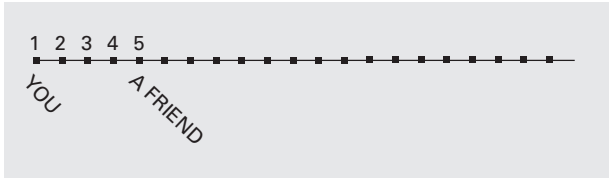

Figure 6.1
A measuring rod directly gives the distance between you and your friend. Image credit: RC.

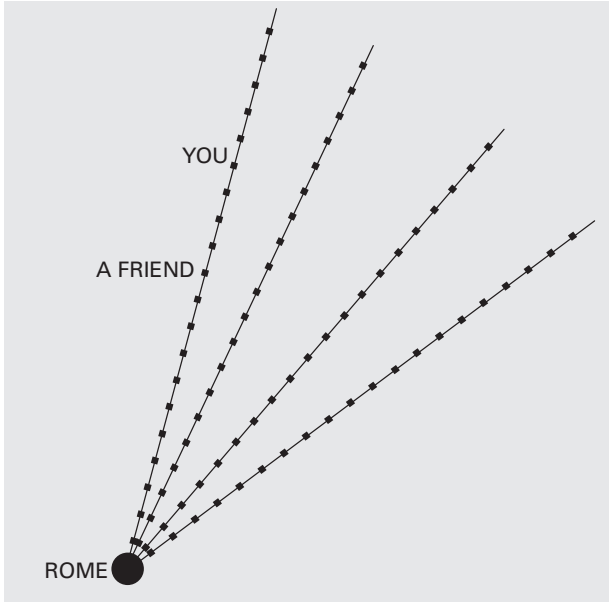

Figure 6.2
A centered web of lines requires that you compute the distance. Image credit: RC.

A stopwatch tells the length of a given temporal interval, of which we mark both time extremities.

`Stopwatch: Input (`T_0`,`T_1`), output: (Interval length)`

A timer marks the most recent end of an interval of given length of which we marked the earlier end.

`Timer: Input: (`T_0`, interval length); output:` T_1

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

The two items are conceptually complementary. In both cases, you fix the older end of the temporal interval; in the case of the stopwatch, you input the most recent end of the temporal interval and the stopwatch outputs its length; in the case of the timer, you input the length and the timer outputs the most recent end of the interval. Think of two uses of a rod. You can align one end with the edge of an item you want to measure and record the reading at the other end (stopwatch-like use). Or you can align one end with the edge of an item you want to cut at a certain given length (timer-like use). In what follows, I shall use, unless otherwise specified, the concept of a timer in a way that does not take into account the differences illustrated in the previous paragraphs. Timers will thus be timers proper, stopwatches, or both.

Finally, note that both timers and clocks include a *counter* (i.e., a physical system that implements a function from periodical processes to natural numbers), and this is an essential feature of both clocks and timers. (Once more, the same holds for rods and measuring webs.)

What powers a clock? A *time-follower* is a device that ticks periodically, and whose ticking periodicity is based on (or sets) some standard time unit. One example of a time follower is an oscillating pendulum. It does not count: it does not output any number. The pendulum just oscillates. Thus, it cannot track time and time intervals by itself (with the sole exception of the single interval that is the temporal unit of its tick or beat). It is not a clock, for it does not indicate time distances from any given t_0, and it is not a timer, for it does not measure time intervals. (Similarly, a rod with marks without any associated numbering would be a sort of space-follower or space-pacer. It divides space into equal intervals but does not count them. Thus, it does not track space and distances by itself.) On the other hand, counting without ticking does not give you either a clock or a timer. Thus, I suggest that if we think of ticking as time marking, we should require that both a clock and a timer incorporate both a counter and a time-marker.

We now can see why we are able to use a timer as a clock, and a clock as a timer. The second case is simple, for a clock already includes a timer as its part. In the first case, we know that we can make a clock out of any timer provided we set the timer. And eventually both a timer and a clock have to incorporate both a ticking mechanism (i.e., a time follower) and a counter. Figure 6.3 shows the tree of the concepts involved up to now.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

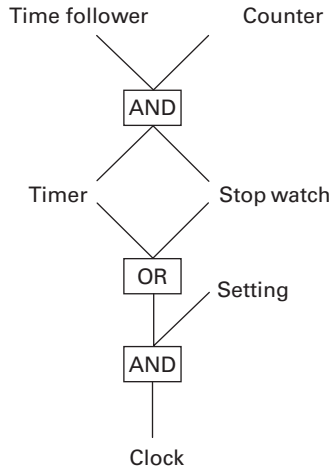

Figure 6.3
The functional architecture of clocks. Image credit: RC.

To count is to input a certain series of items and to output a series of numbers. Time counters do exactly that; they input phases of processes and output the number—a positive integer—corresponding to the present phase. In the case of time, we measure it by counting regular periodic events broadcast by a ticking mechanism that is juxtaposed to what we want to measure. The number of periodic events gives us one temporal measure according to some arbitrary unit of time.

Clocks and Maps

A clock is *not* a map. A clock*face* is. Clocks contain mechanisms that place and keep the user on an event map[6], moment by moment. There is Red Dotting in clocks: the tip of each clock's hand places the user in a certain position in time and strives to keep it at that place so that when certain significant events occur (e.g., the *n*th passage of the Sun on Greenwich meridian on planet Earth, counted from a standard t_0), the hands are in the appropriate position on the clock's face.

So, what kind of map is a clock's face? It turns out that it is very practical to use rotational mechanisms, and this shows on the circular design of many clockfaces. But its hidden logical form is that of a *line*. A couple of compromises are possible, between the necessity of a linear representation

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

of time and the practicality of a spinning mechanism. In both compromise cases, the clockface expands to an endless line. Compare: the Spiraling clock (figure 6.4) wraps the line on a spiral on a cylindrical surface and makes it pass through the tip of the hand for, say, the hours. As time flows, the line moves. The Wheel clock (figure 6.5) is a wheel with a painted clockface, turning at the appropriate rate on a straight line on the ground. In both cases, contact of the line with the wheel or with the hand occurs at the relevant time.

Providing users with just a tiny bit of learning how to (re)use the same map day after day saves clockmakers from the impractical burden of using endless spiraling or straight lines to tell time. The clockface is a *recyclable map* as it gets reused each day. Although it does contain many abstract time representations (representations of intervals of hours, minutes, and

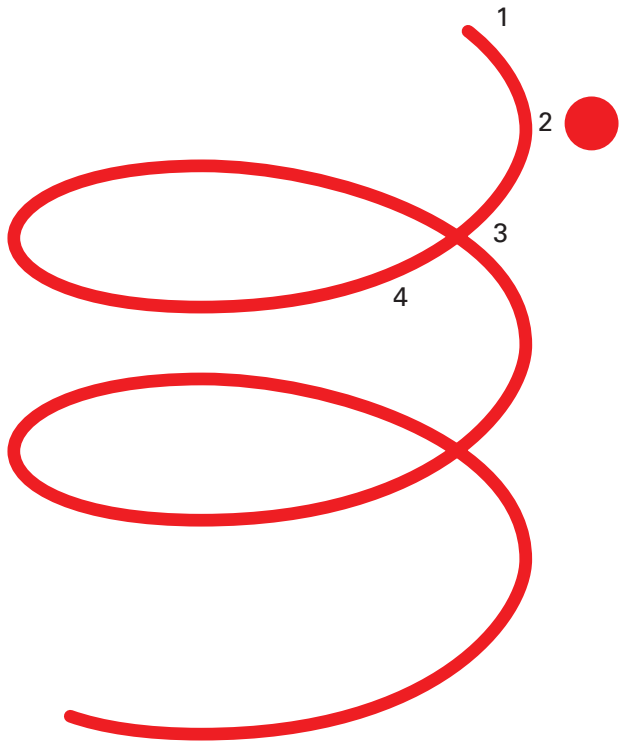

Figure 6.4
An endless helix climbs and passes through the Red Dot of the present. It is two o'clock. You need an infinitely long line to make this clockface usable. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

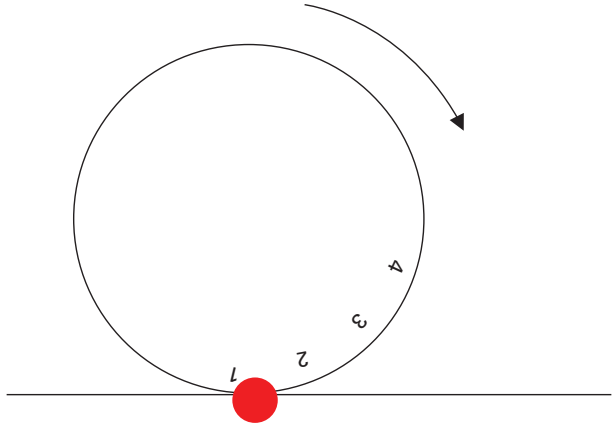

Figure 6.5
As a clockfaced wheel spins on an endless line, the contact point makes for a Red Dot. It is about a quarter past one. Image credit: RC.

seconds), it is the map of an actual territory, that is, a series of events containing one salient point, the passage of the Sun on the meridian today.

The territory of a clock, as I have said, is day after day the whole of a day's temporal content, anchored to a unique daily event, which is the local passage of the Sun on the meridian. This is what clockfaces *represent*. Suppose we lie on Paris's meridian at 9 a.m. The clock hand points to 9, and the clockface has, among others, a mark for noon. That mark, *now*, refers to the impending noon in Paris.

The content of what a clock indicates is best expressed by something like the following clause (assuming clocks could speak): "When I say that it is time t, it is time t" (where "t" stands for the name of a time, expressed in the usual units, dating back to a given zero time). Thus, clocks are like a temporal GPS display, allowing us to navigate time thanks to the event map engraved on their face through a mechanism that strives to keep the Red Dot in the right position on the map.

The Geographic Space of Clockfaces

The nexus between clockfaces and GPS displays is worth exploring. Used on Earth, a clock tells not only the time but also what meridian you are on (modulo time zones). Consider the following inference.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Assume the clock (veridically) tells you that it is noon *now*.

And that, as it happens, now it is noon *in Paris*.

Hence the clock tells you that you are on the Paris meridian or time zone.

To draw this *spatial* conclusion, the only thing you have to know is that it is now noon in Paris.

Accordingly, a clock can *misinform* you about your location. It tells you that it is noon now. And now it is noon in Paris. The clock tells you that you are on the Paris meridian. But you are actually in London! It is 11 a.m., not noon, in London. If at 11 a.m. you look at your clock in London and it tells that it is noon, you may philosophically claim that the clock is not misinforming you, as it is telling what time it is now in Paris. Big Distraction! In your less philosophical moments, you would claim that the clock is plain wrong. Because it tells you that you are on the Paris meridian, and you know you are in London.

Consider a 24-hour clockface. In place of the marks for 3, 6, 9, 12 hours, it has marks for 6, 12, 18, 24 hours. Make it a bit more complex. Instead of the hour hand, a miniature map of the globe as projected from the South Pole occupies the whole of the clock's face. In the projection (an azimuthal equidistant projection), the South Pole is at the center of the face and the North Pole is degenerately represented by the outer circumference. Meridians are represented by straight lines (i.e., rays emanating from the center), and parallels are concentric circles. As we saw in chapter 2, it is as if each meridian was a piece of soft spaghetti held in place by a pin at the North Pole, with the South Pole touching the ground. You unpin the pieces of spaghetti, and they all fall to the ground radially from the South Pole. At the end of each meridian-spaghetti you find an image of the North Pole, which is then spread over a circumference.

The miniature map slowly revolves clockwise around the center, coming full circle after a day. Now, abstract from time zones, with their irregular shapes and quantizing purposes. Have the movement set relative to a standard of time, for instance by setting 00:00 at the Greenwich meridian. Given this particular design, you will get the following:

- Determine where you are, and you'll know your time.
- Or, conversely, determine where in the world it is now the time you read on the clock dial, and you'll know the meridian you are on.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Suppose your position is given by a GPS-based navigator and translated on the revolving miniature map by a Red Dot. The number you'll read at the end of the meridian on which the Red Dot lies is your current time. Take away the map, just leaving the Red Dot and the meridian. This is almost an ordinary clock's hour hand. In an ordinary 24-hour clock, you just cannot see the map behind the meridian. But *the map is virtually there*, and the hand is just a meridian on that map: your current meridian (figure 6.6).

Thus, a clockface is also a map in the sense of literally being a *spatial* map. Under each clockface, there is an unarticulated constituent, a hidden map of the world centered on the South Pole.[7] And if clockfaces are both calendars and maps, then clocks provide Red Dotting to *both*. Through the Red Dot, you can see both a time and a place.

Perfectly working clocks can fail to tell the time. If you are at the South Pole, the Red Dot is at the center of the clock's face. It lies on every meridian.

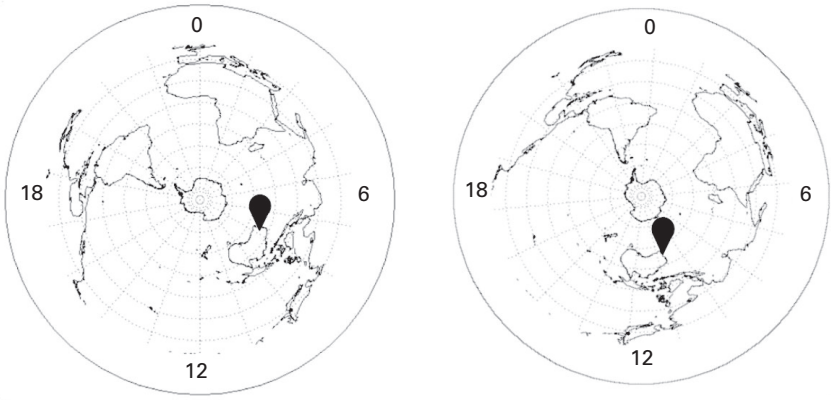

Figure 6.6
An Earth map in azimuthal equidistant projection is mounted on a 24-hour clockface. The hour hand is replaced by an indicator of the current physical location of the clock (the Red Dot, here symbolized by the pin). Centering the projection on the South Pole aligns the clockwise spin's direction of the Earth (as seen from the South Pole) with the traditional clockwise movement of the clock's hand. The picture represents two different times, three hours apart. On the left-hand side, it is about 8 a.m. on (the meridian of) Australia's west coast. On the right-hand side, it is about 11 a.m. The same projection could be mounted on the face of an equatorial sundial set for the southern hemisphere. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

If you are at the North Pole, the Red Dot spreads and occupies the outer circumference where, there too, it lies on every meridian. And fair enough, at the North and the South poles, it does not make much sense to ask for the hour.

(Distraction: We can agree on some convention for that. For instance, as the South Pole is normally accessed via New Zealand, it conforms to NZST.)

Two travelers heading north from different starting points who meet at either pole could only observe the differences between the readings of their clocks. Their clocks at the poles only indicate the direction they are coming from; *they only contain spatial information*.

Consider once more sundials. An equatorial sundial whose stylus coincides with the North (or the South) Pole and is parallel to the Earth's axis just indicates *where* on Earth it is now midnight, by pointing its shadow toward a meridian. On each point on that meridian, it is now midnight (and of course, in some of these places, at midnight the Sun shines).[8]

The Chronogeoscope

Based on the GPS-clockface idea, I designed with Glen Lomax a small piece of educational software, named the "Chronogeoscope" (figure 6.7), that makes it possible to study the conceptual tie between time and place in a clock.[9] Your location on a map of the Earth provides the hour hand of a 24-hour clock. By reading the number on the hour ring corresponding to your location, you know what the solar time of your location is. You are not actually using a clock, but a model of the Earth's rotation, to get an idea of the current time.

We demonstrate that hidden in each watch's face there is an implicit image of the Earth's spin. As per the account presented previously, we divided the time on the watch's face not into 12 hours, but into 24 hours. We then turned the clock's face into the representation of the Earth centered on the South Pole. This way, the Earth's map spins the way the Earth does (seen from above the South Pole, the rotation is clockwise). Then, we set the period of the rotating map's spin to 24 hours. If you are patient enough, you will see that the map gently rotates.

At this point, the key idea is to determine your location on the map. This is done either by entering the location directly in the latitude and longitude boxes or by allowing the software to retrieve your position. Your position

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

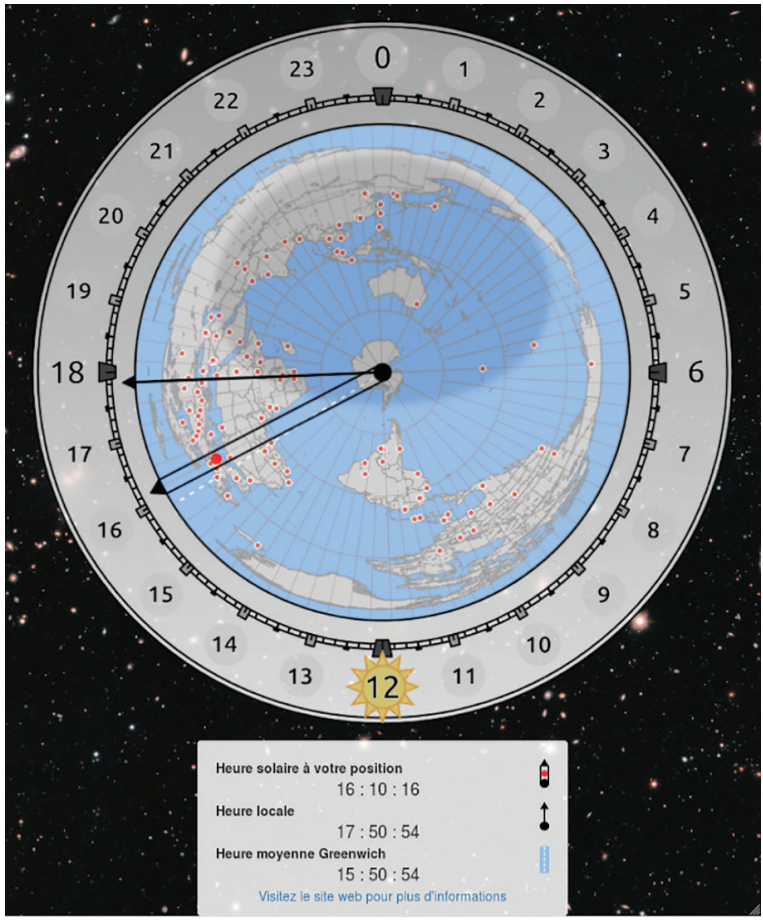

Figure 6.7
The face of the Chronogeoscope. The Red Dot corresponds to your position or to the position you want to know the present solar time of. The thin hand gives the present conventional time for that position. The dotted line represents Greenwich meridian. Image credit: RC and Glen Lomax.

is then indicated by the Red Dot. Your location is, on the hour hand, on a 24 hour–period rotating map of the world. This means that *you can use space to tell time*.

To understand the Chronogeoscope, it is important to understand the way we built it. The Red Dot *does not look at a timepiece that tells the hour but to a GPS that tells the place*. At this point, the Chronogeoscope uses a

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

timer—not a clock—that sets the rotation at the 24-hour period. We added another hand that does take its input from an actual clock. Whereas the Red Dot specified solar time, the thin hand signals the actual, conventional time. For instance, this takes into account daylight savings in some countries, where during the summer you can see that clocks place you to the east of your current location.

The shaded area represents the night on Earth. The Sun is to be found in the direction of noon. The shape of the shadow depends on two factors: day of the year and peculiar features of the azimuthal equidistant projection. In the boreal spring and summer = austral fall and winter (March through September), the shadow engulfs the South Pole and does not touch the North Pole, which in the azimuthal equidistant projection is a degenerate point, as it coincides with the outer perimeter of the map. In the austral spring and summer = boreal fall and winter (September through March), the shadow engulfs the degenerate representation of the North Pole.

As we saw earlier, when used on Earth, a clock tells not only the time but also what meridian you are on (modulo time zones). Take away the map from the Chronogeoscope's face, just leaving the Red Dot and the meridian. This is almost an ordinary clock's hour hand. In an ordinary (24-hour) clock, you do not see the map behind the meridian. But the map is virtually there—and the hand is just a meridian on that map: your meridian.

Many complaints have been addressed to what we may call the northern bias in geographic representation: in most maps, north is up, South is down (Ashworth 2019). The Chronogeoscope is centered on the South pole. This tool displays this orientation because when the South Pole is seen from above, the Earth spins clockwise. As developers, we chose to retain the clockwise-spin correspondence between the display on the Chronogeoscope and how we read ordinary clocks. Had we centered it on the North Pole, the Chronogeoscope would have displayed a counterclockwise rotation. The clockwise constraint and the use of the azimuthal equidistant projection allow us to explore a southern perspective on geographic representation. In the Chronogeoscope, we simply eliminated the up-down pairing and replaced it with the center/periphery pairing: South is at the center, and north is everywhere at the periphery. If you travel to different time zones, you will be able see the difference between local time and the standardized time at your location. (For instance, this difference can be sixty minutes for Paris, as opposed to five minutes for New York.)[10]

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Constellations: A Note on Geographical Clockfaces

Some situated ecological artifacts, or "naturefacts" in Heersmink's (2013) terminology, like the group of circumpolar stars and constellations that one can use as a clock and for other orientation purposes, are maps. Constellations are perspectival accidents. Their visual aspect depends on our viewpoint—the position of the Earth in the Milky Way. But given that our viewpoint changes little relative to the stars composing constellations, constellations end up functioning as visual invariants.

The use of constellations as visual invariants should be contrasted with the use of the *Old Man of the Mountain* or other environmental landmarks (Sorensen 2013). The Old Man of the Mountain was a geological feature in New Hampshire (United States) that, when seen from a certain viewpoint, resembled a bearded man's profile. You could use it to identify the only viewpoint from which you saw the bearded man. This way to engage an environmental landmark is not how constellations are used; we do not navigate in space and then dead reckon where we are because "we see this constellation again." In navigation, constellations are used as a two-dimensional, stable pattern of figures.

The main cognitive advantages of constellations are that they allow us to identify important stars, including through phrases like "Arc to Arcturus, spike to Spica" (Casati 2019a), and that they can be used as if they were clockfaces. As recommended in instructional astronomy books such as (Rey 2006), once you have identified certain stars, you can *then* create a further artifact, the circumpolar clock.

Stars can further be assigned regions on Earth corresponding to the zenithal star (i.e., the star directly above an observer) for that region at a conventional time such as noon in Greenwich on the vernal equinox (Huth 2013). Assigning Earth regions to zenithal stars turns a map of the Earth into a star map, and vice versa (figure 6.8).

Temporal Red Dot Maintenance

As clocks (and the GPS-based navigation devices) show all too well, Red Dot maintenance can be automatized, thus shifted to M4 (a proxy mode). Short of that, we have to resort to M2 (slow and stepwise mode) or M3 (use of external artifacts), at different levels of granularity and precision.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

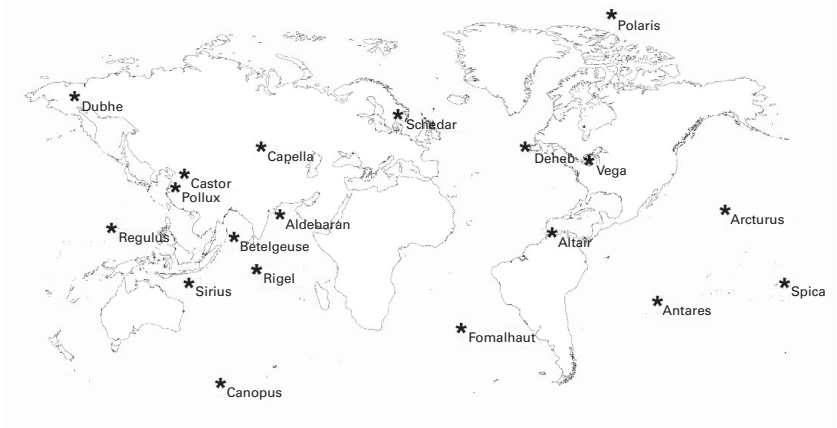

Figure 6.8
Seen from the Earth's center at noon in Greenwich on the spring (or vernal) equinox, stars can be conventionally made to coincide with places where they are at the zenith. Memorizing the superposed star-Earth maps helps viewers find stars given knowledge of the location of other stars. Image credit: RC.

Vintage tear-off calendars required users to tear off the page of each elapsed day; if you forgot to tear off the page once, the calendar would end up with a misrepresentation of the current date. Many operations were involved with tear-off calendars: waiting for the end of the day, tearing off the page of the just-elapsed day, reading the name of the new day, and remembering that one had already torn off the relevant page, so as not to detach one page too many. Some theoretical knowledge is relevant to tear-off calendar interpretation: if you sail around the world, movement across time zones will eventually disrupt the accuracy of the tear-off procedure (as Jules Verne's Phileas Fogg realized when completing his tour in *Around the World in Eighty Days*). Accordingly, epistemic segregation holds for calendric Red Dotting. A webcam on my Parisian tear-off calendar generates a Green Dot image that is not useful, and is potentially misleading, to my antipodal friends in Wellington (New Zealand).

And of course, epistemic segregation holds for clocks as for any other Red Dotting device (figure 6.9).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

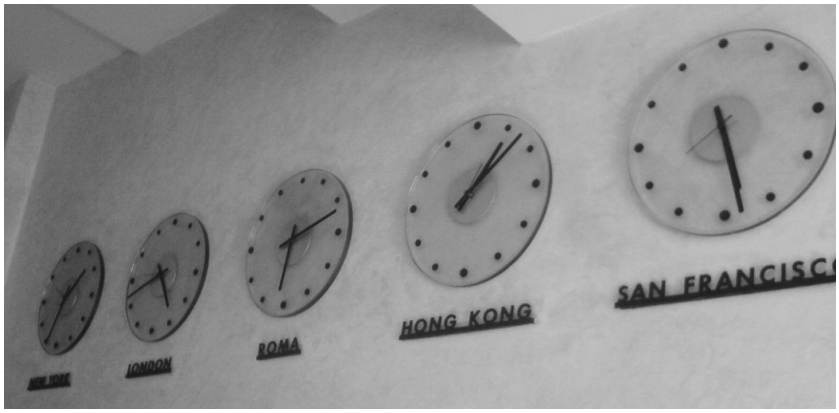

Figure 6.9
Labeling clocks compensates for the lack of epistemic segregation. Before asking the time, find out where you are. Image credit: RC.

Conclusions

Clockfaces present interesting features that depend on the complexity of clocks. We reconstructed the conceptual components of a clock, from ticking mechanisms to counters to clockfaces. In the same way that Borgesian maps were shown to be conceptually prior to other types of maps, the Chronogeoscope is conceptually prior to clocks; a spinning map of the Earth is a hidden representational component of each traditional clockface, which is best understood as a Red Dotting device.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

7 Music Notation as Map

Music Scores

Music notation is a fascinating representational artifact. Standard Music Notation (SMN) as used in Western cultures is graphically complex and, on many accounts, not very user-friendly, as its current format is the result of a sedimentation process that lasted centuries and has involved musicians and music publishers rather than cognitively minded designers. This collective labor has left traces that are difficult to improve upon or reform (S. Rankin 2018). And yet, if you think about the immense variety of sounds, rhythms, styles, instruments, and effects music notation can represent, it is extraordinary that it works at all, in the sense that it is possible to learn how to use it even if it takes years to fully master it. It's a clearly effective system.[1]

Most Western musicians and musically educated people acquire a majority of their repertoire by reading music scores.[2] (Distraction: In many cultures music is transmitted without written notation. Yet, an explanation is still needed for how music notation works for those who do use it.) Learning to read music is a long and time-consuming process. Some crucial conventions must be mastered and implemented following sensorimotor patterns that are specific to the instruments one plays. The sight of a simple notation item such as a musical scale (a-b-c-d-e) "feels" motorily and acoustically different to a flutist and to a pianist, respectively. As every musician and music teacher knows, music notation definitely does not feel naturally generative—it takes thousands of hours of dedicated training to learn to read music proficiently. But once learned, reading music repays the effort an individual put into his or her long training. In this sense, music notation is very similar to writing. In fact, Western music notation systems historically originated from writing,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

starting out as an encoding of musical events by cues superposed on written texts of song lyrics (Strayer 2013). However, my claim in this chapter is that at least *some* aspects of music notation interpretation are naturally generative—indeed, hypergenerative—particularly the aspects that pertain to the understanding of temporal order and that allow navigation of a music score. That is, in some aspects, music scores are a very specific kind of *map*.

Let me immediately qualify this point. (I'm indebted here to Catherine Elgin, who back in 2013 strongly questioned the map claim about music scores.) Music scores are seldom produced with the intention of representing a musical performance. Some are: jazz musician Keith Jarrett's improvisation at his 1975 Köln Concert was first recorded on tape and then transcribed and published as a score (Jarrett [1991] expressed some reservations about the project in his preface to the music.) Interestingly, the transcription has been played by other musicians in public performances and further recorded by Tomasz Trzcinski in his 2006 album *Blue Mountains*.

But Jarrett's performance turned music score is not the general case. Scores are used as maps not by design; they serve other functions because they are *prescriptive*, that is, they guide the trained reader in the execution of a motor task. So be it! We still have to explain how we can use them in certain circumstances.

Turn that page, *now*! Consider this very specific music activity: as a trained musician, I sometimes go onstage to turn pages for pianists at rehearsals and public performances. What is *my* kind of cognitive engagement with the score? I have to know moment by moment what the pianist is currently playing by listening to it, and I have to find on the score the item that *currently* corresponds to the played section of the music. For me, as a page turner, the score is a representation of the performed music. And I use the score as a map to orient my behavior ("the moment is coming for me to turn the page, quick, prepare, now, act!"). Some scores that are available online as read-along for musical performances display a thin moving line that fulfills the role of the Red Dot. For the collaborative page turner, music scores conform to the map principle of Red Dotting; maps are items such that one can (mentally, or physically, by pointing with your finger) place a Red Dot on them. You can place functional Red Dots on music scores, so they are maps. This specifies also what the territory is for music scores. Music scores are maps when they map the temporal unfolding of concrete musical events that is, music performances. Musical events are

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

their territory. Other map uses of music scores occur when a pianist loses her thread in a performance of chamber music and quickly glances at the score to see "where" other musicians are. Or, when a student works on a music dictation and repeatedly listens to a certain passage, trying to make sense of what she wrote on the staff. If Red Dotting is possible for music scores, this is because the latter also satisfy the other structural aspects of maps, in particular the same-dimensionality constraint. This means that we can expect a score to be translatable as a formal score in the same way in which a map can be translated into a formal map.

The translation of Western music notation into a formal map requires some specific work. The temporal aspect of music is coded by semi-arbitrary symbols (they are semi-arbitrary in the sense that some of their meanings can be derived systematically). For instance, once you grasp the relationship between the symbol expressing an eighth note and the symbol expressing a sixteenth note, you can predict the meaning of the symbol expressing a thirty-second note. The temporal length of a note is not coded by the spatial length of the symbol. But this is no more problematic than the fact that some geographic properties are not represented by regions of space on ordinary maps. For instance, being a two-lane highway does not need to be represented by a double line; it can be expressed by a color or by a particular hatching.

Figure 7.1 is an example of how a music score can be translated into a formal map.

The standardly notated melody is unpacked as pitches, each of which is assigned a hatching (or a color) and located durations (intervals on a continuous line). For instance, the first G (=horizontal hatches) occurs before the first A (=45° hatches), and the second lasts half the duration of the first.

Pitch is not a territory but an abstract space (prepare for possible Distractions here). We choose to represent pitch in formal music scores with hatches and not with the organization of spatial regions, as we do in formal maps for expressing properties. On the other hand, events in time are a territory, and we use one spatial dimension to represent their temporal unfolding. It should be noted that while Standard Music Notation represents durations conventionally and pitch analogically, formal scores represent durations analogically, and pitch conventionally.

Figure 7.2 shows how the formal score of figure 7.1 is unpacked in a series of score-layers, each presenting the locations and durations of one

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

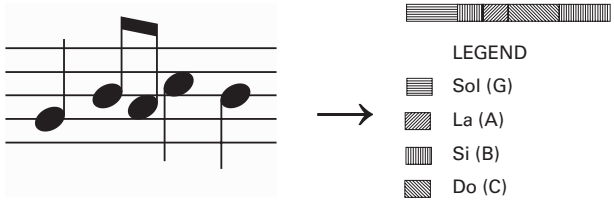

Figure 7.1
Translating a score fragment (left hand) into a formal score (right hand), through a legend. Pitches are represented by hatching: relative temporal durations by relative spatial lengths. Image credit: RC.

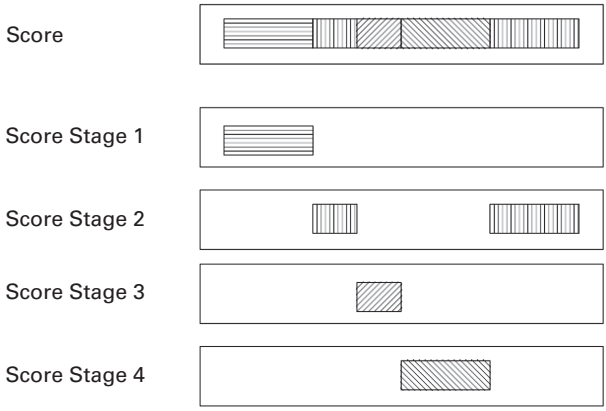

Figure 7.2
Layers (stages) of the formal music score of figure 7.1, here superposed in coordination along the temporal line. Image credit: RC.

pitch at a time, interleaved with locations and durations where the pitch is not present. Each layer is true if at the temporal location denoted by a colored area we find the note corresponding to that color, and only if no such note is found at the location denoted by the blank portions of the layer. The truth of the score is then a function of the truth of each layer.

In theory, different players could be assigned the task of playing a single pitch each, and their coordinated performance, made possible by the temporal superposition of each layer, could in principle produce the melody, albeit awkwardly. (Needless to say, we are not proposing formal scores as substitutes for standard notation. They are just providing the formal structure of an aspect of standard notation.)

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Music notation does not work like formal scores (although it could). In the following sections, we'll look into a characterization that is closer to the surface features of music notation, while still enhancing music notation's map-like features. A new, intermediate notion will prove crucial, the notion of a *location raster*. But we should reject the general non sequitur according to which "From the fact that music notation can be expressed as a formal map, it does not follow that music notation is a map." Not everything in a music score is map-like, but something is, and it is important to highlight it.

(This discussion invites a philosophical Distraction. If scores are maps, can we conclude that there is a territory they map in situations other than that of the page turner? Are we led to endorse some platonistic theory of music objects? Are we committed to such an engaging metaphysical option?)

The next sections explore some aspects of the conventions related to time representation in music scores. They introduce a syntactic characterization of a fragment of SMN and discuss some cognitive consequences of principles that govern the syntax.

The Temporal Fragment

SMN is a notation that primarily represents the evolution of pitch in time. In its present historical form, pitches are represented as locations on a five-line staff. Here we shall abstract from pitch and present a characterization of the time representation dimension only. This is meant to provide the semantics of a *fragment* of the complete set of SMN notation symbols. In formal semantics approaches to natural language, it is common practice to provide fragment accounts of complex phenomena (e.g., the modal or the temporal fragment of natural language). As such, the music fragment abstracts not only from the pitch dimension but also from many notational peculiarities and redundancies, in particular. In the following, we shall start with a simplified version of the fragment and then add a few principles that capture many of the idiosyncrasies of SMN.

The simplified fragment includes a set of primitive symbols.

Notes, such as , ,

and a functional expression, the tie mark

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Notes are aligned on what we shall call a *raster*. A tempo signature (e.g., 3/4, 6/8, etc.) is provided to ensure that notes aligned in the appropriate way on a time-signed raster constitute a music *score*.

The notion of a raster is central for the possibility of using SMN (as well as many other systems, as we shall see) to represent events in time. A raster is an abstract organization of locations on the score, characterized in terms of some of its spatially relevant features. A raster R is an ordered set of linearly arranged discrete *spatial* locations:

$s_1, s_2, \ldots, s_n$

such that s_1 is the leftmost element on the raster, and each s_{m+1} location is to the immediate right of its predecessor s_m.

A *connected segment* of the raster is a set of positions of the raster such that if s_m is the leftmost element of the set and s_n is its rightmost element, all s_i such that $m \leq i \leq n$ are part of the set.

An *initial connected segment* is any connected segment that includes the first (leftmost) location of the raster, s_1.

Given a raster R, a *score* on R is the result of filling every location in a given connected initial segment of R by a note symbol.

Figuratively, a raster is a kind of sequence of locations s_n (here designated by placeholders, "x"):

x x x x x x x ...
$s_1, s_2, s_3, s_4, s_5, s_6, s_7, \ldots$

It gets filled location by location, starting from the leftmost location.

♪x x x x x x x x x ...

Let us pause and take a look at what happens here. First, some observations on the potentially distracting "spatial" structure of the raster and the shape of symbols. Left-to-right orientation (customary in SMN) is actually immaterial to the organization of the raster; the notation could be otherwise oriented. But nothing of importance hinges on its being oriented the way current practice has it. Spacewise, the raster could unfold in a spiral; again, nothing depends on this—so let the raster be visually arranged on a straight line. Notes could have triangular or pentagonal shapes, or different colors to differentiate them; nothing depends on this, hence let us use symbols that are sufficiently similar to those used in current practice.

Second, we need a characterization of the raster structure because we want it to represent time. What can a raster represent? The representation

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

provided by the raster is mereotopological only (i.e., it is meant to preserve mereotopological properties only). The simple semantics of the raster is the following:

> If a location s_m is to the left of a location s_n, then the symbol placed on s_m denotes an event that precedes the event denoted by the symbol placed on s_n.

Given that the only relevant properties are order properties, the raster is not to be construed as implying a *pace*. In particular, and importantly, the referents of the locations on the raster are not requested to be equally spaced in time.

Third, and crucially, the raster structure is *discrete*. Locations in a raster are atomic. There is no such thing as a half location on a raster, or as the leftmost part of a location.

These features make rasters (figure 7.3) rather different from other types of map scaffoldings such as coordinate grids. Coordinate grids are *paced* (equally spaced according to some metrics such as distance or angle). They carry information about *both* the underlying spatial structure (possibly empty space) *and* the spatial size and structure of the items placed on the grid. The spatial equivalent of a raster would be a rack of shelves on which descriptions of geographic items or symbols representing those items would be placed. The shelves would not be sensitive to the size of the items referred from those descriptions or symbols.

Raster structures are little noted but important representational structures. They underlie other types of notation. Lists, and in particular verbal descriptions of itineraries, are rastered (figure 7.4).

Alphabetic writing systems are rastered too, as they are, in part, a representation of the unfolding of phonetic events in time. In some particularly transparent writing systems, if a letter symbol is placed to the left

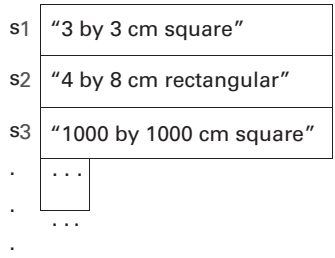

Figure 7.3
A rastered map.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

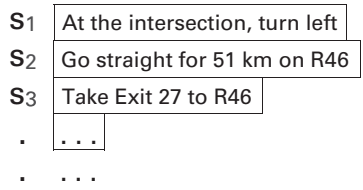

Figure 7.4
Itinerary descriptions are rastered.

(or the right) of another symbol, then the referent of the former must be pronounced before (or after, respectively) the referent of the latter symbol. The raster for writing in most Western alphabets can be filled in by replacing empty position by letters:

d x x x x x x x x x . . .
d o x x x x x x x x . . .
d o g x x x x x x x . . .
d o g s_a n d _ m i c e . . . [3]

Here too, the preserved structure is purely mereotopological, size-insensitive. A character count (including spaces) of a written text does not tell how long it will take to pronounce the text, either in its entirety or with regard to its different parts. As it happens, the pronounced length of a piece of text depends on the particular assignments of lengths to the referents of the letters as well as many other factors. The same vowels in given contexts take more time to pronounce than in others. The spacing of locations on the raster does not make one pronounce the letters at any pace. In both this case and in the case of music notation, the raster is necessary to represent the unfolding in time; the raster is discrete, and its only relevant spatial structure is mereotopological structure. (A generalization to other types of writing systems is possible. Syllabic, ideographic systems, at some level, are raster-like.)

Note that, on this account, written linguistic items are maps but are not maps of the *content* of the items. If Anna utters the words "Mary runs," a written sentence like "Mary runs" is a map of the linguistic sounds emitted by Anna, but it is not a map of the situation the uttered sentence describes. This echoes the case of pictures (see chapter 3), where we understood that from the fact that a certain representational vehicle, such as a picture, is a map, it does not follow that the vehicle is a map of what it represents.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Molecular Expressions

An important feature of SMN is that it allows for *molecular* expressions. Using primitive symbols opportunely placed on a raster, one builds molecular expressions by *ties* connecting symbols at adjacent locations.

If note *a* occupies position s_m and note *b* occupies position s_{m+1}, say that *a* is *adjacent* to *b*.

A *dyadic molecular expression* is constituted by two adjacent notes connected by a tie.

Now, stipulate that ties can take as their argument both atomic and molecular expressions. Thus, a molecular expression can be constituted by a note *a* at s_m linked by a tie to a molecular expression in turn composed by note *b* at s_{m+1} and note *c* at s_{m+2}.

However, we do not have any reason to distinguish the molecular expression constituted by a note *a* at s_m linked by a tie to a molecular expression composed by note *b* at s_{m+1} and note *c* at s_{m+2}, from the molecular expression constituted by a molecular expression composed by a note *a* at s_m linked by a note *b* at s_{m+1}, and note *c* at s_{m+2}.

Visually:

is identical to

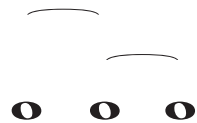

and both are identical to:

Hence, by convention, a n-ary *molecular expression* is constituted of pairwise adjacent n notes connected by n-1 ties.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Call a *concatenation* any set of notes and/or molecular expressions such that it occupies a connected area of the raster. A *music formula* is any concatenation of notes and/or molecular expressions. Thus, music formulae can overlap (share constituents).

Not all music formulae of the temporal fragment are well-formed from the point of view of SMN. SMN gives preeminence to temporal units and introduces a number of notational abbreviations. To appreciate the rationale for the idiosyncrasies of SMN, let us now turn to the semantics of our fragment.

Semantics of the Temporal Fragment

The SMN lexicon is simple. Notes represent events (sounds or silences). SMN distinguishes between pauses and "sounding notes." The distinction is irrelevant for the purposes of the present discussion and will not be leveraged in the discussion of the temporal fragment. We shall talk about the referents of notes generically as "events," without distinguishing between silences or sounds.

The temporal fragment represents the unfolding of events in time. The temporal asymmetry of this unfolding is mapped into a reading direction. As we observed, the only intuitive temporal significance of spatial features of the raster is ordinal: the event denoted by a note at a location s_m precedes any event denoted by a note at a location s_{m+n}. The raster is only used for writing down formulae; it does not represent or imply pace or a beat. In the fragment, information about temporal measure is completely deferred to notes.

Thus, notes come in a structured, open-ended lexicon. They represent events and assign them a duration. Current convention states that starting from the unit note, which is assigned an arbitrary time value, notes can express any of the $1/2^n$ values of the unit note.[4] Starting from 𝅝, the unit note, the lexicon includes 𝅗𝅥, one half of the unit note, and then ♩, ♪, . . . one quarter, one eighth, . . . one $1/2^n$th . . . of the unit note.

Molecular Expressions

A molecular expression composed of notes or of molecular expressions *a* and *b*, and a tie between them denotes a single *whole* (temporally connected) *event* whose duration is the sum of the durations of *a* and *b*, that is: 𝅗𝅥⁀♩. In the example, the denoted event's length is 3/4 of a unit (1/2 + 1/4).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Music *formulae* are concatenations of notes and/or music expressions. They denote sequences of events of various durations.

The following are two examples of music formulae:

, and

We have seen that a *score* on R is the result of filling every location in a connected initial segment of R by a note. Some notes may be connected with ties, thereby giving rise to molecular expressions. *Scores* in SMN come with a meter indication. *Paced* music scores are music scores preceded by a meter indication.

 or 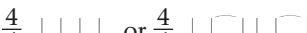

A music score denotes a *music piece*, that is, the set of events that are the referents of atomic or molecular expressions, and that unfold in time as the score indicates.

Map principles govern the interpretation of scores. We have seen in chapter 2 that maps are not idle, not glutty, not gappy, and not ambiguous. Likewise, on top of avoiding ambiguity (an atomic music expression does not denote two distinct events) and gluttyness (two distinct music expressions do not denote a single event), the temporal fragment is constrained by some sort of fullness (no gaps) and nonoverlap (no idle parts). A music formula has *no temporal gaps*. Between the events denoted by any two adjacent expressions in a formula, there is no room for a third event. The temporal fragment is further constrained by *nonoverlap*. The temporal endpoint of the event denoted by the leftmost note in a couple of adjacent notes is always earlier than the temporal starting point of the event denoted by the rightmost note in the couple.

Taken together, fullness and nonoverlap endow the temporal fragment with the possibility of using notes and molecular expressions to systematically follow a rhythmic pattern. If the notation was not full, one would not know the starting point of a given event once the previous event is terminated. And if the notation allowed for overlap, one may have multiple events running at the same time, with no systematic way to determine the relative start and endpoint of any of them.

The Temporal Fragment and SMN

In SMN, conventions are used to render some molecular expressions; for instance, one ½ note followed by two dots is definitionally equivalent to

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

a molecular expression composed of a half-note, a fourth-note, and an eighth-note, as in

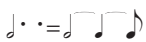

In general, each dot written to the right of a note or of another dot corresponds to a length that is half the length of the note or the other dot. Dots are explained away by the definition (but see below for exceptions to the generalized use of the abbreviation). More interestingly, SMN uses explicit indications of *bars* (i.e., temporal units for a rhythm). These are *redundant* in the temporal fragment, as the temporal unit is specified by the time signature. However, explicit bars facilitate music reading, and possibly music writing. Related to bar units are some rules of well-formedness for SMN. In fact, not every formula of our simplified temporal fragment is a formula of SMN. Indeed, SMN adds to the elements of the temporal fragment that have so far been described what we may call the Bar Limit Principle.

Bar Limit Principle: Notes (i.e., atomic expressions) cannot straddle bars.

This is an example of a well-formed formula (WFF) according to the Bar Limit Principle: . And this is an example of a non-WFF:

The Bar Limit Principle does not forbid suitable *molecular* expressions to straddle bars. Thus, although the two following formulas are semantically equivalent, only the first is a WFF in SMN: ;

Further conventions forbid certain ways to concatenate notes in a molecular expression whenever the bar limit is *not* trespassed. For instance, or are not allowed if the events denoted by these molecular expressions are entirely contained within the temporal unit of a single bar. The very same expressions may however be allowed if the denoted events straddle bars at suitable points.

We observed that in general dots are definitionally explained away by the use of ties. However, definitional equivalence does not allow for uniform substitution. Dotted notes cannot cross a bar limit. The following, for instance, is not allowed (the vertical line marks the end of the first bar):

In this example, the first illegal dot is assumed to be related to the second note in the score.

This completes the characterization of the most salient features of the temporal fragment of SMN. The characterization shows that distinct representational time aspects are assigned to the raster structure (a purely

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

mereotopological, discrete, and ordered structure) and to the temporal values of notes. The raster structure is used to express temporal order. Notes are used to temporally measure events. Some principles (fullness and nonoverlap) regulate the interplay between order and length of the represented events. Finally, meter indications subtly interact with the rules for using molecular expressions through the Bar Limit Principle.

Specificity of Fullness and Nonoverlap

Fullness and nonoverlap are not general requirements of the semantics of natural languages. "Cleopatra run and Nefertiti run" does not require the two runs (i.e., the semantic values of the two event sentences) to be adjacent and disjoint events. It may be asked whether the items of SMN obey a stronger principle, Fodor's Picture Principle, according to which if *a* is a picture of *b*, then parts of *a* are pictures of parts of *b* (see chapter 3). Fodor's Picture Principle is meant to tell linguistic representations from iconic representations. (As such, the principle does not distinguish well between pictures and diagrams, but this is not important for our present purposes.) Music scores indeed comply with the principle but only modulo the raster structure that individuates legitimate parts of a score. The only parts of a score that have meaning are either atomic or molecular expressions. "Areas" of a score that straddle the right half of a note and the left half of the subsequent note are not meaningful. Indeed, they are not syntactically individuated parts. However, once the restriction to formulae is accepted, the parthood structure of formulae mimics the parthood structure of the events that are their values.

According to the present characterizations, SMN is akin to a formalized language. It is formal also insofar as it is a human artifact, and as it obeys stipulative rules.

I underscored what appear to be interesting analogies with *writing*. Writing too, as it is used to represent the temporal unfolding of spoken language, is a kind of formalized language whose rules are, of course, proprietary and different from those for music notation. Writing also complies with Fodor's Picture Principle, modulo the restrictions to parthood imposed by the raster structure. Music notation and writing are thus iconic, *discrete* formal systems. Thus characterized, they both come close to *maps*, whose logical structures are underscored by the exploitation of some topological regularities. This is why music notation and writing can be used as maps.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Conclusions

This chapter provided a syntactic and semantic characterization of the temporal fragment of Standard Music Notation. We stressed the role played by the theoretical notion of a raster. The raster is an "invisible" grid of discrete locations that can host symbols according to some map rules and principles (fullness and nonoverlap). The grid itself has a simple topological structure. Importantly, items in the grid are atomic locations, whereas subatomic locations would not be vehicles for representation. Modulo this structure, raster-based notations are map-like, accept Red Dots, and can be used to navigate music performances.

Raster-based iconic representations—where elements of the raster are atomic—also include written language as a means of representing the temporal unfolding of spoken language. These forms of iconic representation are halfway between (paradigmatic) diagrams and (paradigmatic) pictures; they are akin to maps and to formal languages such as the ones used in logic.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

8 Organizers as Maps

The Alpine Sequence

Italian children used to learn a trick to remember the somewhat contrived names of the subranges that constitute the Alps and their positions in the overall Alpine range. The trick consists of pronouncing the first syllable of each subrange's name one after the other, which by a lucky coincidence create a sentence with some sense. This trick is a mnemonic device. For Italian children, the Alpine sequence to remember alpine subranges is:

Ma co' gran pena le reca giù ("However, with great pain she brings them down").

The names of the subranges are:

Marittime-**Co**zie-**Gra**ie-**Pen**nine-**Le**pontine-**Re**tiche-**Ca**rniche-**Giu**lie

from the Thyrrenian sea to the Adriatic sea, following the Alpine range from west to east. The west to east direction is coded in most maps as a left to right direction, thus facilitating the association. If the names of the subranges were inverted, the children learning the Alpine sequence would have to manage an added layer of complexity.

The written or spoken Alpine sequence is a map. It has parts, and these parts refer to parts of the territory, respecting the systematicity constraint. The Alpine sequence can be used for orientation, and a Red Dot can be placed and maintained on it.

The map-like nature and use of the Alpine sequence is directly confirmed by a small debate about alternate sequences. Schoolteachers and pupils can choose between:

"Ma co' gran pena le reca**no** giù"/"Ma con gran pena le **reti** cala**no** giù"
("However, with great pain they bring them down"/ "However, with great pain they bring down the nets").

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

These examples add "no" for the **No**riche Alps, but in an incorrect position, as they are actually west of the Carniche Alps.

Another possibility is:

"Ma co' gran pena le **reti non** calò giù"

("However, with great pain she did not bring down the nets").

This possibility adds the Noriche Alps and locates them in the correct position.

Schoolteachers and pupils seem to be aware that there is a necessity to correctly represent the order of the mountain subranges. They are also aware that this spatial representation is quantized: for whichever item is considered as spatially meaningful in the phrase, its parts are not thereby spatially meaningful. The syllable "Ma" refers to the Maritime Alps, but the "M" in "Ma" has no spatial significance, and neither has a spatial significance the trans-syllabic item "ac," a proper part of "m**a c**on gran . . .".

Each Alpine sequence is learned by heart. Learning by heart has cognitive costs. But these costs are made bearable by splitting them between first packaging the order information in an easily learnable sequence and then memorizing the names of the ranges that are opened by the first syllable in the Alpine sequence. It is difficult to remember the sequence "Marittime-Cozie-Graie-Pennine-Lepontine-Retiche-Carniche-Giulie," but it is easy to remember the short, meaningful, and vivid phrase "Ma con gran pena le reca giù," and to subsequently use relevant syllables as prompts for the full names of each subrange.

In English, a comparable mnemonic device is the geological sequence, which reads "Egyptian Camels Ordinarily Sit Down Carefully, Perhaps Their Joints Creek Tremendously Quietly," referring to the periods Ediacarian, Cambrian, Ordivician, Silurian, Devonian, Carboniferous, Permian, Triassic, Jurassic, Cretaceous, Tertiary, and Quaternary. It is not a map as it does not support Red Dotting, except if you visit places such as the Grand Canyon, where the geological pile is laid bare in the rock layers. In many cultures and languages, a number of word sequences are learned by heart. A beautiful example is the sequence of star names organized along the rising points of stars on the horizon in Polynesian and Micronesian cultures. The Hawaiian version, developed by Nainoa Thompson, is based on the teachings of master navigator Mau Piailug and includes for the east-to-north rising points:

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

(east) Puana, Kauluakoko, Humu, Hōkū'ula, Hōkūle'a, Nānāhope, Nānāmua, Keoe, Hawaiki, Hōkūlei, Nā Hiku, Iwakeli'i, Hōkūpa'a (north).

(In Western names, these star names translate sequentially as Procyon, Betelgeuse, Altair, Aldebaran, Arcturus, Pollux, Castor, Vega/Lyra, Deneb, Capella, the Big Dipper, Cassiopeia, Polaris.)

Navigators learn sequences of star names when they are ashore and then use the sequence to identify stars as they rise and later on symmetrically set relative to Polaris. In the absence of written communication, tokens are laid on the sand on a circular line—a star compass—they are named, and the names in the sequence are learned by heart (Gladwin 1974; Lewis 1972; Kyselka 1987).[1]

Sequences of numerals, of names of the days, of the days of the week, of the months of the year, of the letters of the alphabet, and names of musical notes, among other more complex structures, such as the solutions of a multiplication table, are the result of dedicated teaching and memorizing practices. With the exception of the list of the star compass, not one of the sequences listed here is a map in the sense that none actually maps a physical territory; or at the very least, it will take some arguing to consider, for example, numbers and days as territories themselves. However, what can sequences of numerals and of names of the days teach us about maps and their cognitive lives?

Consider numerals. Names such as "one," "two," and "three" are not introduced independently of the introduction of many other names, such as "seven" or "ten." Numerals up to "three" are learned in a nonrandom way, hence disregarding the fact that their semantics is available to the learner thanks to subitization (i.e., the immediate, precise grasp of small numerosities, what is considered initial knowledge); they are learned as ordered *parts* of a nonrandom sequence of other numerals, disregarding the fact that for these other numerals the learner has no available semantics. More broadly, primitive numerals are learned independently of a semantic assignment. You do not even try to introduce "seven" alone. You spend a lot of time instructing children to learn numerals in the appropriate order.[2] This effort is comparable to the one of learning a rhyme. Why invest time and energy in a rhyme that includes names for numbers the kid has no "initial" access to?

It is not numerals *per se* that are a maps or map-like but certain *sequences of numerals*. Sequences contain at least two items and, in this case, two or

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

more adjacent numerals. Sequences are not just orderings; they are map-like insofar as they have a mereotopological structure. The sequence "3, 4, 5, 6" contains, among others, two adjacent parts, say the portion "3–4" and the portion "5." Some of its parts contain other parts ("3, 4, 5" contains "3–4"). In Casati (2014), I discussed the idea put forward by Spelke and Tsivkin (2001) that numeral sequences provide a representational *bridge* (as defined in chapter 1) between two core M1 cognitive capacities, that is, subitization and rough estimation of large numerosities. My claim is that numerals bring indeed precision but do not refine approximate representations; grasping some of the properties of numeral sequences depends rather on memorizing them and seeing that they have a map-like structure. The account has the advantage of explaining why we use numerals such as "seven." "Seven" is a primitive numeral—it is not the result of any combination of numerals. If we had positional numerals in base three, so that our "four" were pronounced "one-one," and seven "two-one," then it would be immediately possible to think of the combinatorial feature as *the* relevant link between small cardinalities and not so small cardinalities. But as far as I can ascertain, no natural language has a numeral system in base three, and "seven" is invariably translated into other languages as a primitive numeral. Hence the combinatorial feature of language, which for Spelke and Tsivkin is the computational edge of numerals, is not necessary: a primitive "large" numeral such as "seven" can denote sharply a collection that none of the two nonlinguistic systems of subitization and rough estimation recognizes with the desired precision.

Organizer sequences must be learned by heart because they *cannot be naturally generated*. There is nothing in the name "three" that tells you what the name of three's successor will be, no matter what your competence with numbers is. Compositional structure applies to larger numerals, but its application presupposes that one has learned the names of small numerals and of other significant numerals (e.g., "twenty," "a thousand") in the appropriate sequence.

Notice the parallel with names for the days and other word sequences. They are learned as rhymes, are introduced in an ordered, wholesale way; are noncompositional; and once learned, their sequence (for instance the series of names of the week: "Monday–Tuesday . . .") can be literally seen (or figuratively, as one should better say "heard") to have map properties; they can be situated (one cannot jump from the referent of "Thursday" to the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

referent of "Saturday" without passing through the referent of "Friday"); and finally the span "Monday–Tuesday–Wednesday" is seen (or heard) to be larger than the span "Monday–Tuesday." This also applies to the names of the letters of the alphabet, to the Alpine sequence we introduced at the beginning of this chapter, and to the sequence in the star compass. In all these cases, a representational epiphany takes place.

Word sequences are *map-like* in the sense that they have clearly individuated parts that are ordered, and that we invest in learning them. Whether they are full-blown maps or not depends on whether or not there is a territory they map. By drawing a comparison to maps, so far, only the *structural properties* of numerals are assumed to be understood; nothing has been said about the intended semantics of numerals. But of course, once the map features of the sequence of numerals are understood, the available, automatic semantics for small numerals can be entered. However, the learner need not assign *any* particular representation to numerals larger than the small numerals. It is enough that the learner be able to use the map, in the sense of being able to operate with it. Thus, we may not have a *concept* of "seven" (and even less of "fifty-three") in the sense in which we have a concept of "three."

I'd like to prevent two possible misinterpretations of the main hypothesis defended here.

(i) The thesis is not that in understanding numerals *all* elements of map use and cognition are retrieved or activated in full, but only that some relevant aspects are taken into account.

(ii) There is a certain holistic[3] flavor attached to the idea that numerals are (somewhat) represented wholesale. However, the thesis is different from the holistic thesis that to master the concept of three one must have mastered the concepts of one, of two, or even of seven (in the sense in which it is said that to master the concept of a teacher one must master the concept of a pupil). The latter is a semantic thesis that is not needed here. Instead, the idea is that we have a tacit *syntactic* representation (the sequence S), not a tacit semantic representation.[4]

What are the consequences of the proposed account? According to the thesis, numerals and names for the days of the week are *small "artificial languages" encroaching into natural language*. The thesis is compatible with the claim that written numerals may have predated spoken numerals.[5]

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

For, if numerals are not part of language and are incorporated elements, then where do they come from in the first place? The answer may be that they got incorporated into spoken language from written language or from iconic representations. If not written language, numerals could have just come from gestures. In either case, we *learned to pronounce a map*: primitive numerals are a map that you can pronounce.[6] This would be an interesting case in which *spoken language* is parasitic on something else—on non-spoken language or on other public representational instruments.[7]

The account is also compatible with the fact that some cultures use names of body parts as numerals. In point of fact, body parts are not assigned randomly to numbers (i.e., it is not the case that once you use your left-hand thumb to denote 1, your left-hand index-finger denotes 5). Hence, the structured sequence of all body parts used as numerals can be used as a map in the proposed sense.

The account proposed here is that the sequence of primitive numerals is represented syntactically as a single unit and that this large unit is partly interpreted in the way a map would be. Some consequences of the thesis have been hinted at in the previous sections: the learning effort for numerals is explained in a straightforward way since it takes time for one to memorize the long sequence; the account is compatible with known aspects of map cognition; it is compatible with the possibility that written numerals have been introduced before spoken numerals; and it is compatible with the fact that in some cultures numerals are represented by body parts. Furthermore, the map interpretation of the account supports both ordinal and cardinal construals of numerals, and makes a plausible place for the invention of "seven." Finally, the account is general insofar as it also covers other non "strictly numeric" examples, such as names of the days of the week, of months, of fingers, of letters of the alphabet, of musical notes, and of subranges in the Alpine sequence. Numerals are in no way special.

Contra Spelke and Tsivkin, numerals are *not* what make it possible to improve number cognition. They are *all there is* as the next steps in number cognition. Put otherwise, there is no number cognition on top of the first, subitizing step, and there is no number cognition for large numbers in the sense that the approximate system cannot be "made" precise by the use of numerals. The number 57 is invisible to unassisted cognition. In this deflationary sense the present account resonates with Wittgenstein's views:

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

> Our language can be seen as an ancient city: a maze of little streets and squares, of old and new houses, and of houses with additions from various periods; and this surrounded by a multitude of new boroughs with straight regular streets and uniform houses. (Wittgenstein 1953, §18)

However, the sequence of primitive numerals does not fall in either category. It is neither a clumsy part of the old city nor a neat part of the suburbs. It is like a fallen deity, a beautiful crystal placed in the very center of our village.

Organizers and Organizations

Numerals are map-like but they are not fully maps as they have no territory to map to begin with, unless one subscribes to some form of idealized but concrete existence for numbers. Nor are some other organizer sequences maps (e.g., names of music notes). Much as it need not be argued that pieces of lands and sequences of events (candidate paradigm territories) exist and beg to be mapped, it will require some arguing to claim that numbers and music notes exist. As announced in chapter 1, I'm happy to be noncommittal about abstract objects; I can even go a step further and assume they do not exist, for the sake of the argument (Casati and Soldati 1995). So here we have map-like things that are not really maps.

In certain contexts, numeral sequences do work as maps. Think of the numeral line atop the entrance of vintage elevators (figure 8.1), with stenciled numerals for each floor.

In the example on the left of figure 8.1, Red Dot dynamics is active (i.e., the indicator is situated in the elevator's cabin, which is on the fourth floor). In the example on the right, Green Dot logic holds; from your location on the ground floor, you see the elevator moving toward the twenty-first floor. (Remember: Green Dot devices track the position of an object on a map; Red Dot devices track the position of the map on itself.) The series of floors in a building is a particular type of territory: it is a sequence of items with an internal structure. Call it an *organization*. Organizers are maps of *objects in sequences* or of *organizations of objects*. There are many organizations in the sense defined here, typically artifactual items. Lines of people waiting for a service are organizations. So are chess tokens at the beginning of a game, hotel rooms, the keys of a piano,[8] rows in a concert hall, the bus stops of the M11 line, gates at an airport, book chapters, and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 8.1
Elevator buttons and floor indication. Interestingly (figure on the next page), the maps can be horizontal, thus abstracting from the verticality of the territory (stacked floors). Image credit: RC.

so on and so forth. Humans find it convenient to organize iterative items, and find it convenient to map them with the sequence of numerals. Most likely, organizations and organizers coevolved to a certain extent.

Inventories also use the numeral sequence (e.g., complex artifacts like cars or sailboats are assembled from several thousand items, each with a sequentially numbered code), but the sequence itself is not used as a map, as the assembly is not an organization.

Organizations are artifacts that are *fit for mappability*. Given their structure, organizations internalize the semantic features that make them mappable at a very low cognitive cost because their elements can be directly and

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 8.1
(continued)

progressively numbered. Organizations possess an order that is easy to sort out. Floor numbers in a building are not assigned randomly or from top to bottom. The generation of maps in the case of organizations is simple: just produce a sequence of numerals.[9] Finally, organizations make the task of drawing Borgesian maps onto them particularly easy. Just label the items progressively!

At this point, we can say something about metro maps, which have often been presented as a special cases of maps, that is, topological maps. They are said to preserve topology and to abstract from metric and other nontopological properties, such as orientation (do not use them for navigating aboveground, as stations that are only a few minutes' walk apart are

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

connected by underground lines that will take you on a long ride). I think a case can be made for a slightly bolder claim. We can consider a metro line as an organization: a spatially organized set of metro stations, connected by tunnels, that as any other organization is fit for mappability. And indeed, each metro line is typically mapped by its own one-dimensional map, one you can typically see printed inside the train wagons of that line. It is a map: the mereotopology constraint is preserved (the map region representing stations CD is for instance contained in the map region representing stations ABCD), the same-dimensionality constraint is preserved (a 1D map represents a 1D territory), and you can Red Dot your position on the map. Mereotopological inferences can be drawn from the map (figure 8.2): going from Châtelet to Censier-Daubenton will take more time than going from Châtelet to Jussieu.

But at the city level, where multiple lines exist and are represented, these inferences are not generalizable. At the city level, the global representation

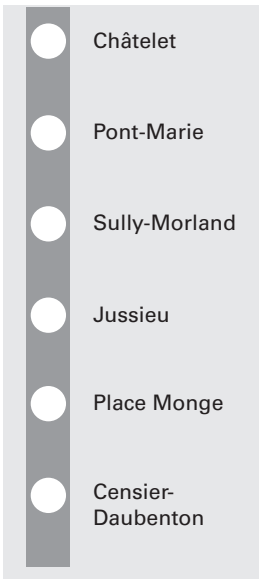

Figure 8.2
Distances between the dots representing the metro stations accommodate the placement of typographical items and do not reflect the distances in the territory. Some inferences about relative lengths are still possible: the metro journey from Châtelet to Censier-Daubenton will take more time than going from Châtelet to Jussieu, as Jussieu is placed between the other two stations. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

of the whole set of metro lines should be construed not as a map but rather as a *diagram* for coordinating with each other all the individual metro line one-dimensional maps, and hence for coordinating inferences about itineraries that only take place on metro lines. Metro lines do not normally intersect; they normally pass at different depths over each other—there are no traffic lights underground. Metro trains do not move from line 6 to line 7. Points on the diagram suggesting a connection between line 6 and line 7 indicate that at some station of line 7 one can get off the train and follow a pedestrian path to a station of line 6 and vice-versa; the two stations receive the same name ("Place d'Italie") but are in fact (in the norm) different stations (figure 8.3).

Non-Map Organizers

As said, the norm in numeral sequences and many other organizers is that they do not have territories (organizations) they map. Can they be used *as if* they were maps? Their syntax is conducive to games of make-believe with them, games of taking them as maps. Much depends on the accomplishment of the learning phase, at the end of which we have memorized the sequence (e.g., "Do-Re-Mi-Fa-Sol-La-Si"). Once the macro-syntactic item is

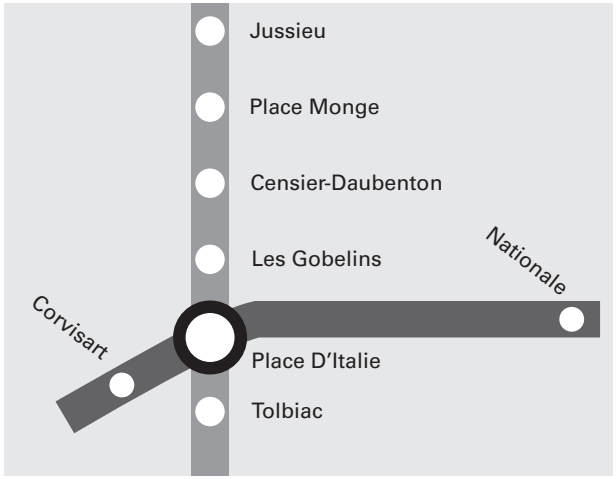

Figure 8.3
Examining the intersection of line 6 and 7 at Place d'Italie, we get no information about which runs above which other. Image credit: RC.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

generated and available, we know that whatever the referent of "La," it will be between the referent of "Sol" and the referent of "Si"; that "Mi-Fa-Sol" will be part of "Do-Re-Mi-Fa-Sol"; that we can generate "Sol-Fa-Mi" once we have stabilized the "Mi-Fa-Sol" vehicle; that if we are "located" at the "Sol" symbol we have a long way to go to get to the "Do" symbol, and so on. Remember: all these examples are *pronounced* sequences. What we learn by rote is the pronounced sequence, and the syntactic operations just described are operations on that sequence. The representational vehicle is a certain spoken event.

The map-typical make-believe games we play with organizer sequences are just that: an engagement with adjacency and parthood relations, an engagement with directions, and Red Dotting. All this can be done by *pretending* that there is a concrete territory that is mapped by the organizer.

By contrast, a thermometer does *not* contain a map. The ordinal representation of degrees in a thermometer from lower to higher bears only a superficial resemblance with other ordinal syntaxes that are found in maps (e.g., numbers for floors in an elevator). There is no territory that is "mapped" by the graduation bar of the thermometer; temperatures do not have a spatial organization (if they happen to have it, this is only by accident). There just is an ordinal structure in temperatures: temperatures are not a territory. Fifteen degrees is between 16° and 14°, but there is no sense in which the "region" denoted by the sequence "14°–15°" is "part" of the "region" denoted by the sequence "13°–14°–15°–16°–17°." Compare: in an elevator panel, the region containing the numerals "13," "14," and "15" includes the region containing the numerals "13–14," and the referent of the former region contains the referent of the latter. One could say that the thermometer graduation syntax is not map-like either. In a thermometer graduation, there are no degree regions, only points, and if it made sense to single out a region 13°, that region would not correspond to an actual range of temperatures that the thermometer would measure.

Queues: Dynamic, Ephemeral Borgesian Maps

In some countries queuing is an obligation or a point of honor; in others, any means would seem legitimate to avoid it. Part of the folklore are the symmetrical frustrations of an Italian person in London and a British person in Rome. The stereotypical Italian in London is frustrated because he understands that he won't be allowed to use one of the many stratagems to

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

avoid waiting in line. Among the strategems: pretend to ask for innocent information directly at the counter, pretend to know someone at the head of the queue, go straight to the head of the queue as if you were going back there a second time, or simply apply the golden rule to go unnoticed set out by John Le Carré: imagine that you are the owner of the place. The stereotypical Brit in Rome does not know how to get to the only open counter: behind *which* person should he stand in line? Even beyond cultural differences, the queue is a complicated institution. The British even have coined a word specifically for this action, whereas an American will use the composite phrase "wait in line." How people choose to stand in line appears to be deeply rooted in social cognition, defining something that is almost an ethics of queuing (Efrat-Treister et al. 2019), which seems to resist the possibility of introducing economic incentives and disincentives to induce orderly queuing. But the deepest explanation of the peculiarities of queuing behavior might be one of a perceptual nature: the queue is an ancient institution that lives on its ability to make visible to an individual who approached a good or service dispensed at a single point how he or she should behave compared to all the people who came before the individual. The individual shifts from M2 to M1 her computations for the strange human activity of waiting.

Queuing is an *institution*, in the sense that queuing does not impose itself as an obvious solution to a problem of coordination between people. Rather, queuing presupposes the recognition of some shared rules. If an extraterrestrial entered a post office and saw people in line, they would immediately understand that this is some form of social behavior. The probability of people randomly distributed in a room being lined up is near zero. The extraterrestrial would think, "If they are in this configuration, there must be a reason." The line is certainly very salient perceptually, and this is not a secondary factor of the success of the line institution. It is immediately *evident* that since *everyone is* doing a certain thing, we are asked to tune in to the collective behavior, or at least to inquire about it.

The basic rules of queuing can be made explicit. Queues are made by people who must wait to access a good or service. Thus:

First rule: *whoever arrives first gets in contact with the place providing the service.*

Second rule: *whoever does not arrive first places themselves behind the last person in line.* (In particular, if there is only one person in line, this will also be the last person in the queue, so the latecomer will be queuing behind them.)

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

A mathematician could simplify the first two rules by saying, "Whoever arrives places themselves behind the last person in line"; if there is no one in line, this rule becomes the first rule and, if there is someone, the second rule. But we prefer to leave two rules for convenience of presentation.

Third rule: *you are not served until the one or the one behind whom you are standing in line is served.*

From the third and first rules, it follows that whoever arrives first is served first: if you are not served first despite having arrived first, either the first rule (someone has jumped the queue and got in front of everyone) or the third rule (someone is served from anywhere in the queue) has been transgressed. From the second rule, it follows that you cannot sneak in a place in the queue other than the last one. This last theorem is modulated by a further principle:

Fourth rule: *whoever does not arrive first can only be placed in last place, unless there is agreement between all the people queuing that the last come can move to the place he or she wants in line.*

From the fourth rule, it follows that it is not enough to ask a person in line permission for stand ahead of her when you join the queue: all the people behind the person you ask would be harmed. On the other hand, the rules (and their consequences) do not rule out that you can exchange your place with someone else in line. The exchange, seen as a private transaction, is tolerable because nothing changes for people who are not directly affected by the transaction.

The queue, in a metaphoric sense, is a machine for *turning invisible time into visible space* since it represents and shows a temporal priority (abstract and invisible in itself) as a spatial priority, which literally jumps out at you perceptually. Think about the fact that people who line up all look in the same direction. (Imagine the following set up: with a group of friends, line up at a counter at your local Post Office so that all of you have your backs to the counter, and watch the behavior of the customers who come in after you. What is *the last* place in the queue?) This transformation of time into space has clear cognitive advantages when it comes to the *ease of drawing conclusions from a visual examination of the line*. As we have seen, a first immediate result of this organization is to suggest to anyone who approaches the counter the existence of a structured social situation. The presence of the queue says: "No anarchy here, there is a social event to be part of if you want to get to the counter." Conversely, being a member of

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

a queue also means *performing* a role for all those who see the queue; by showing that you are in the queue, you indicate the rules of the game to those who want to play along. Finally, there are subjective cognitive advantages to queuing: the fact that you are in line allows you to keep an eye on only one customer at a time, the one who is immediately before you in line.

There are certainly other ways to organize the waiting time for several people in front of a counter. You can walk into a waiting room and ask, "Who is the last one?" But then, what about the next customers? If a person comes in right after you and does *not* ask who is last in line, you have perhaps an entitlement to say out loud, "I am the last one." But what happens if yet another person comes in and also does not ask the question? Then, it would seem that the rule, in this case, is to leave it to the person who has entered after us to inform the newcomer. But what if he or she doesn't say a word? To take the floor and say, "Look, *you* come after that person (who, in turn, comes after me)" seems too strong, the typical unsolicited intervention of a fussy, impatient troublemaker. A *queue manager* would be then needed, but there is no shared social rule that allows this manager to emerge automatically, a self-proclamation would be seen as potentially unjust, and then the hypothetical manager could no longer perform her function once it is her turn at the counter (unless she is *really* prepared to play the role of the troublemaker!) The queue tacitly solves all these problems: the order does not have to be declared verbally but is inferred by whoever approaches the queue from a summary examination of the way the waiting people are arranged. And, of course, if one slips in front of everyone, reproaches will be voiced.[10]

The basic rules are *constitutive* of queuing. If one could change their place without the explicit approval of the others in line (in particular the individual with whom she swapped her place) there would be no line to begin with. If the first in line was not served, or if we do not wait until the person in front of us is served to try to access the service ourselves, the line would not be a useful tool. There are also less basic rules, and as such more flexible ones, which depend largely on the context. For instance:

Fifth rule: *you cannot hold a spot for another person unless you give up the right to your place, or unless the other person arrived at the same time as the person holding her place*.

The context will suggest when to apply the rule. Standards may be high for lines to buy theater tickets. If, on the other hand, an entire family has

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

to check in at a hotel, it is relatively normal for only one member to stay in line, and this is all the more acceptable if the member in question does not arrive half an hour before the rest of the family. Conversely, it is not acceptable that Alice keeps the place for her other fifty schoolmates, who all show up only when Alice is a few steps away from the counter. Among the reasons for dissatisfaction that could be cited by those who are waiting behind Alice is that this behavior does not allow a fair assessment of the *size* of the queue. "It seemed to me that I was to enter a line of four people, and now I found myself fifty-fourth in line!" The fifth rule is in tune with the idea that the queue is a *perceptual signal* allowing people to draw conclusions about what to do. If one were to keep a place for others in a way not regulated by the rule, one would violate one of the basic rules and create "perceptual disorder" or noise that prevents those approaching the line from understanding what to do. Queues shift the waiting procedure from a complex M3 sets of activities to quick M1 perceptual examination.

The Paradox of Queuing

If the queue has some cognitive advantages (e.g., its immediate readability, the large number of inferences we can draw from a summary examination of the queue, and sparing continuous assessment of the status of your priority), this does not mean that our intuitions about queuing are always clear. In particular, it is not clear for those who queue whether queuing is a useful thing. That there is no clarity is evident in cases such as boarding at an airport gate, where a traveler can assess with good approximation how many people are queuing. No passengers other than those we see at the gate are expected to arrive, except marginally the usual latecomers who are chased by loudspeakers throughout the airport. People line up for boarding, and often *everyone is in the queue*, even if those who are last in line, and precisely *because* they are last, have *no reason* to do so.

The opacity of intuitions crystallizes in a real paradox:

(a) If you are the last person of the queue, *every second spent in the queue is wasted* (Schwartz 1975). If one knew that one would definitely be the last person queuing, as in the boarding queue, one could do something else.

(b) If you are in any position other than last, you are not wasting your time. In fact, you are gaining time compared to whoever comes next in line. If you arrived after them, you would spend more time queuing.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

From these two facts it follows that in order for you not to waste your time in line, you must hope that for the same amount of time spent in line someone would come and line up after you! But this is problematic. The time spent in line, with or without someone *behind* you, is the same. Think about the following, slightly surreal, conversation where Person A says, "I wish I wasn't the last one in this line." Person B responds, "In the sense that you would like to exchange places with someone in front?" And Person A insists, "No, in the sense that I would like someone to enter the queue after me."

The solution to the paradox is simple; it is our intuitions that are confused. If you are certain you are the last in line, the best thing for you is really to go away and do something else. The explanation of the difficulty we have with our intuitions must take into account the delicate balance between the many factors that determine the behavior in line. The cognitive advantage of the queue, based on its evident structure with simple visual inspection, comes at a cost. For the queue to be *perceptually* salient, it is necessary for each participant to do her part (i.e., perform the role of the person in line), even if she knows or fears she is wasting her time. But in this case, being last in line can involve significant costs, since it wastes *all the* time spent in line and on top of that it *frustrates* the performative work of queuing since there is no one after the last one who benefits from the perception of an ordered queue. The marginal benefit for the last in line is that it conforms to a cognitive script, one of being an informant for those who come after (even if the information has no practical consequences).

For this reason, an effective substitute to the queue are numbered tickets, which relieve those who wait from the constraint of having to perform for the benefit of those who arrive later. A numbered ticket also allows people in line to do something else while the queue is running. Numbered tickets are a real creative progress over staying in line.

The complexity of the queue, expressed by a number of implicit rules, derives from the need to keep in balance several cognitive and noncognitive factors. These factors include the fairness of access to a service (whoever arrives first is served first) and the readability of the queue (i.e., its ability to signal to newcomers that there is no other way to access the counter). We have also seen that the cost of readability can be too high, for example, when the last person in line stays last. And, in this case, an improvement in the "design" of the queue is the attribution of numbered tickets.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

But certainly the queue is a social institution, and is surrounded by powerful taboos that are difficult to negotiate.[11] Here is an observation I made in the field. At the exit end of Newark International Airport (New Jersey), in the middle of a tiring night of canceled flights and missed connections, a long queue forms for the cabs. At a certain point an elderly person arrives slowly pushing a trolley loaded with luggage, approaches the head of the queue as if there was no one waiting, and asks for something to the first person in the queue, who is almost about to take the first unoccupied cab. Then, she asks something to the second person. Finally, the elderly person joins the third person and after a brief conversation, she gets into the car with him. Those of us who are far back at the end of the queue mumble. We are soon after told that the lady was looking for someone who had a destination close to hers and that in exchange for the possibility to get into the cab with him or her, *she paid for the ride*.

Now there is something in the lady's behavior that seems "not quite right," but what is it? She hasn't "stolen" anything. She's paid her fare. Of course, the person has also made a profit. A profit is made, in the sense that even though the elderly person paid for the ride (which she would had done anyway), she gets to arrive at her destination long before she would have had she queued, and gets to avoid the inconvenience of the queue itself. And yet, no one seems to lose: not the first two customers in line, who still took the cab when they were supposed to take it, and not the third person in line, who hit the jackpot—thanks to the elderly person, he gets a free ride. For all those who are between the third person in the queue and the place where the lady should have found her position when she arrived, nothing has changed. Those who arrived after the elderly person also benefit because they have one less place in the queue ahead of them (the lady has merged her place with the third in the queue). Perhaps, it is the taxi driver community that loses because of this merge (two customers have become one), but it is doubtful that the slight resentment (or at least the perplexity) among those of us who witnessed this scene is due to a form of empathy or identification with the category of cab drivers.

The lady is violating the paradigm, the script of queuing, introducing a deviant behavior compared to the prototype. But which aspect of the script forbids her to do what she did? The rules do not prohibit the *purchase* of a place in line. Should we then imagine that there are some additional rules that we have not considered? A rule like the following one?

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Sixth rule (?): *you cannot sell your place in line.*

Let's ask ourselves what may be the reasons that make us resentful or perplexed. Maybe the elderly person does not pay the cost that we "pay" in terms of time by standing in line. But then, why did we not go there ourselves and ask the first in line? Or, the elderly person arrived last and left first, and therefore her behavior is unfair to the person she bought her position from, who leaves when she leaves but had to wait before leaving with her. However, the elderly person paid and the person "sold" his place in line, and besides, why should we, from the back of the line, sympathize with the first person in line?

The Queue as a Perceptual Institution

At end of the day, the rules of the queue serve to keep the queue alive in situations where it is useful not so much to *use it* but to *recognize it at first sight* as an institution aimed at regulating access to resources. This primacy of perceptual functionality explains the preference for queuing over asking who is the last person in line (in the latter case, there is no perceptual salience); this explains why the rules of queuing are what they are: they keep the queue alive and make it aggregate and flow in an orderly manner. We can also discuss the visual *relevance of* the rules of the queue; visual relevance explains why replacing multiple queues to multiple counters with a single queue is an acceptable innovation, as it reduces the visual disorder of multiple queues in contexts where it is possible to apply this approach. Visual relevance also explains why you cannot keep a place for others in the queue unless you both give up your place and *make* it clear that you are giving it up; or why changes in the queue are not well tolerated.

The visually tidy queue makes an *unruly* behavior immediately salient, as the latter generates visual disorder. The primary perceptual function of the queue also explains why it is difficult to innovate, where innovations tend to create visual clutter. It also explains why it is not possible to buy or sell a place in a queue except, exceptionally, in a nongeneralized way: when multiple actors are involved, the inevitable stock exchange–like "shouting" would reduce the perceptual effectiveness of the queue by drowning it in an all-too-loud background noise. Alterations to the informational structure of the queue are inferentially costly; they take you from visually inspecting the queue to making inferences from it in M2 state (e.g., the elderly person that bought the place in the queue; you have to plug in knowledge,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

hypotheses as in "What did the elderly person mean to do?") The paradox of the last person in line (who wastes all of her time in line if she is really the last one) indirectly testifies to the fact that the queue is useful especially in situations where it is necessary to *see* that there is a queue; finally, replacing the queue with a series of numbered tags that indicate a given order reduces the costs for those who have a low priority number, as it allows them to do things other than standing in line. Innovations are accepted where the queue is by its nature invisible and is a pure order, as is the case with waiting on the telephone for a call center operator (where the queue is called an "acquired priority"). Tacit queues are manipulable: impatient customers can be managed, but only silently, without letting the others know.

Queues move, albeit slowly at times. If they do not move, or if newcomers feel the queue is not moving, they would lose their function. Thus, they are dynamic maps, and Borgesian ones at that, as they represent the position in time of individuals trying to access a service by the position in space of those very same individuals. The organizer here coincides with the organization.

Conclusions

This chapter analyzed organizers, i.e., ordered sequences of representational items that display map-like behavior, in the presence of abstract or actual territories, the latter being organizations. Organizations are particularly structured territories, that are fit for mappability. Organizers are oftentimes the object of dedicated learning, as in the case of numerals and names for temporal series or of stars in star compasses. Sometimes they emerge relatively spontaneously as ways to regulate collective behavior, as in the case of queues, which work as dynamic Borgesian maps.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

9 Concluding Remarks: The Cognitive Life of Maps

The complex spatial network of Venice[1] is somewhat hard to master and represent, as many nonnatives discover when visiting the town. Repeated visits and strolls through the city provide some comfort because individuals improve their mental representation of and their ability to navigate Venice. And yet, a sense of disorientation (Fernández-Velasco and Casati 2020) remains as the individual moves through the city. Commercial maps of Venice are typically hard to decipher, and many are not even accurate at a level that matters for pedestrian navigation. The irregular shape of the city, its sheer size, and the tortuous paths through it certainly contribute to the difficulty. A complex numbering address system creates further confusion (Denis 2018). It may be that the main difficulty is related to the geographic idiosyncrasies of Venice. The spatial "atoms" or basic objects of Venice are islands, but they do not correspond to the mental spatial "atoms" or basic objects of the pedestrian. In fact, the islands are barely visible to the pedestrian. Whoever takes a gondola tour has the impression of visiting a different city than the one she visited walking. The ontology of Venice is not in line with the granularity of the ontology used by a pedestrian visiting it.

Excluding the mainland portion of the municipality, "prototypical" downtown Venice is an archipelago composed of more than one hundred island-like entities. "Real" islands—like Sant'Erasmo or San Servolo—lie outside the downtown area. Downtown Venice is a tight set of would-be islands. They are would-be only insofar as this portion of Venice is connected to the mainland through a 4 km-long bridge (*Ponte della Libertà*), carrying car and railway traffic, and insofar as most of the component islands are connected to the other islands by one or more bridges.

In the past, canals held a different status. They were the passageways of the upper class, while the tiny streets on the ground were used by house

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

personnel and workers. This explains why majestic entrances face canals whereas pedestrians can hardly find where a building is. Presently, Venice's canals are mostly used for the transportation of goods, public transportation, and sightseeing tours on gondolas. In the past, Canal Grande was the equivalent of any town's Main Street. Now, people stroll and meet in tiny streets between the train station and the Piazza San Marco. This partially explains why on the maps of Venice islands rather than blocks are traditionally dominating: in the past, canals were simply more important than streets.

An island is roughly some portion of land that can be *circumnavigated*. Bridges do not hinder circumnavigation, provided your boat is sufficiently small. This representation of the city is, however, the sailor's point of view. What about the pedestrian's viewpoint? When you're walking in a city, your movement is limited by *blocks*. Blocks are enclosed areas that are standing in your way when you try to get somewhere. Usually, they consist of a group of buildings or just a single building. They can be very large and complex, enclosing private gardens, industrial areas, ports, airports, or highways. If a large block contains also a lot of dead-end streets, as is the case in the historic center of Venice, it's like a trap or a maze. In the case of irregular, large, and complex blocks, we are rarely aware of their shapes and exact location. We know that they are there somewhere, and we try to avoid them. Where are these blocks, and what shapes do they have?

A block is a portion of urban area you can *walk* around, or (to coin a new word) that can be *circumpaced*. Not any walkable path around a portion of urban area defines a block. Blocks are defined in terms of minimally enclosed areas that can be circumpaced but cannot be crossed. A square ("campo" in local Venetian terminology) can be circumpaced, but it can also be crossed. Hence, a closed walk on a square does not include a block. Some borderline cases of blocks should be excluded, such as those defined by paths around a statue on a pedestal, or around a tree. When relatively trivial exceptions are dealt with, Venice has a number of blocks that fit the definition. Some of them are included in islands; some of them straddle islands. Interestingly, none of the Venetian blocks coincide with an island. Compared to mainland cities in Italy, some blocks are anomalous and idiosyncratic to Venice and maybe some other water towns.

How can these different ontologies be captured in maps? Are the mapmaker's choices influencing how the city is perceived? To visualize the complexity

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

of the overlapping webs of blocks and islands, we created an alternate representation of Venice. We first digitized a tourist map of downtown Venice, on which we then drew our own maps. The key to the block map was to omit dead-end streets (including dead-end squares and dead-end courtyards). In the island map of Venice, we kept dead-end canals as not to lose information about canal banks and bridges that are lining the blocks. Also, we assumed the Ponte della Libertà to be severed, so as to exclude the mainland from our calculations. The tourist map that we used as a starting point, as well as the additional maps used for cross-checking, were not fully accurate for our purpose. Particularly, the areas with a lot of *sotoporteghi* (covered passages), such as the residential area behind Venice's train station, were difficult to make sense of. Therefore, we decided to go onsite to gather information. Between January 2014 and March 2017, we walked all the streets of downtown Venice and checked whether they are walkable or dead-ended.

Under these assumptions and based on the information we gathered, we drew two large-scale maps of downtown Venice: figure 9.1 represents only islands, and figure 9.2 represents only blocks. The block map shows pedestrian's Venice in that it depicts areas that cannot be crossed and have to be

Figure 9.1
A rather recognizable map of downtown Venice. Pedestrian paths, including bridges, are not traced. Islands only are visible on this map. Image credit: RC and Magda Stanová.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 9.2
The cutout shows the same area as figure 9.1. It is a much less recognizable map of the center of Venice. Water paths are not traced. Blocks only are visible on this map. Image credit: RC and Magda Stanová.

circumpaced. The pedestrian's map differs from the island map, so much so that many familiar features of the city become unrecognizable.

When the two maps are superposed, the result of the graphical interpolation is the map in figure 9.3. In this map, light blue indicates water; pink indicates walkable areas; gray indicates buildings, publicly non-accessible areas (e.g., private gardens), and dead-end entities; and white indicates bridges. Interestingly, buildings, private places, and dead ends as well as bridges are by-products of the superposition of the two maps. Once blocks and islands as primitives are given, some entities are represented in compositional fashion. Water is the complement of islands, and blocks are the complements of walkable spaces. Bridges are superpositions of the complements of blocks (i.e., walkable areas) and the complements of islands (i.e., water); and buildings, private spaces, and dead ends are the superposition of islands and blocks. There is one exception: the four small blocks on the top of the Rialto Bridge represented in light blue are also buildings (although very small ones) over water. Casati, Stanová, and Roisin (2017) document some of the conceptually challenging choices that were made concerning a number of limit cases.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Figure 9.3
The superposition of the island map and the block map generates a somewhat canonical map of Venice. Image credit: RC and Magda Stanová.

Considered purely topologically, there is a perfect duality of islands and blocks. Figure 9.4 represents an idealized model where each island straddles several blocks and, conversely, each block straddles several islands. But even as we can point to a geometric duality between islands and blocks, they certainly differ in other aspects. Islands are paradigmatic examples of geographic objects: an island is always self-connected (i.e., two separate pieces of land cannot be considered one island) and maximal (i.e., an island is always a whole piece of land, not just a part of it). Islands consist of land, or of land and water, when there is a lake inside an island; islands never consist of water only (but: are icebergs islands?) Blocks, on the other hand, can consist of land, water, or both. We could say that, while islands are generally mono-composite, blocks are either mono-composite or multi-composite. Also, blocks are rarely maximal pieces of land or maximal pieces of water bodies. Combined with a human bias toward maximal connected uniform regions,[2] these factors disadvantage blocks in geographic representations. In other words, there is a sort of *island advantage* in geography. (Even the districts of Venice [*Sestieri*] are defined as groups of islands and not as groups of blocks). This island advantage adds nontopological aspects into geographic representation, which are not necessarily in favor of pedestrians'

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

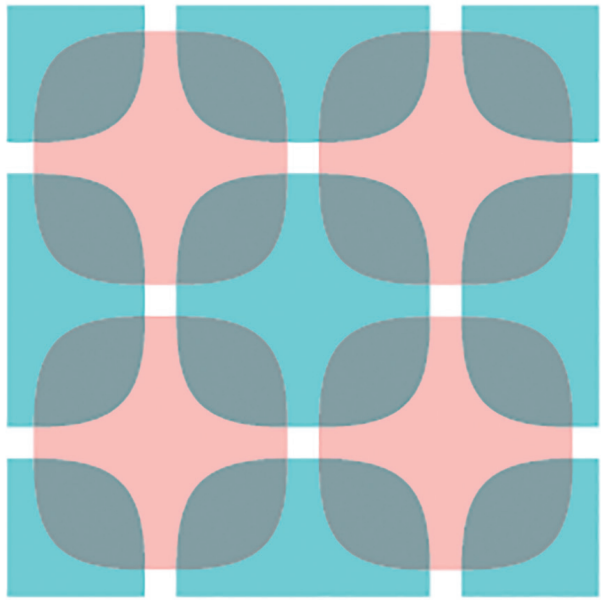

Figure 9.4
Duality of islands and blocks. From the perspective of pedestrian navigation, only some parts of islands matter; the rest is noise that makes more relevant information less visible. Image credit: RC and Magda Stanová.

navigation. The discrepancy between the two forms is all the more evident when you consider that we counted 124 islands and 1154 blocks.

Compared to Venice in 1729 (Casati, Stanová, and Roisin 2017), the number of islands has slightly decreased while the number of blocks has increased dramatically. More than 90 percent of these extra blocks are internal blocks (street only + streets and banks + bank only). Back then, the majority of the internal blocks used to have banks; now, there are many more internal blocks that don't have a bank. So, if you are strolling through Venice and feel that you are encountering water all the time, keep in mind that you would have come across even more water in the past.

The block representation (figure 9.2) can be challenged because of its very high degree of abstraction. As islands are not represented, it is perhaps difficult to locate oneself on the map. However, the block representation does not need to compete with the island-cum-block representation (figure 9.3) on all accounts. In particular, from the block representation, it

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

is possible to draw a number of interesting inferences about "walkability" that are harder to retrieve from the island-cum-block map. We immediately see:

- shortest paths between two points
- sectors of the city that connect to the rest at only a couple of points (e.g., a large area at the north of the city in the Cannaregio district)
- detour options
- wider walkable open spaces (squares)
- that it is easy to circulate within one area, but it will be hard to move to other areas
- that the path from the entry point to Venice (Ponte della Libertà) and to the eastern parts of the city is not straightforward
- that not-walkable areas ("traps") are generally larger than just water areas

To enhance the block map for orientation purposes, we can indicate the water/earth interfaces. In figure 9.5. blocks are depicted only by displaying components that surround them—streets, banks, and bridges. Islands are not fully represented; only island contours that are relevant to pedestrians appear on the map.

Figure 9.5
Bridges (red) and banks (blue) are made visible on the block representation, raising its ceiling. Image credit: RC and Magda Stanová.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Besides the advantages of the block map listed previously, the map in figure 9.5 makes immediately visible the path with fewer bridges, which is valuable information if you are carrying a load since almost all Venetian bridges include stairs.

We are not proposing that the block map or any of its enhanced versions should provide a substitute for a map that combines the island-cum-block representation. After all, the duality of island and blocks is what makes Venice so distinct from other cities. We can see in the maps possible ingredients of a pedestrian's spatial representation of an environment that, in spite of its complexity, displays a number of important constraints, and we can assess the usability of a map that abstracts from other, more conventional elements. If the effects of the island advantage are minimized, some information that is important for pedestrian navigation suddenly stands out. And, although they may not be obvious at first, the block representation has reasons to be commended: a purely island representation is simply not usable for a pedestrian, and an island-cum-block representation is less readable when it comes to seeing connections between salient points of the city. A block map enhanced by information about water/earth interfaces makes, then, a basis for a pedestrian map.

Do Maps Create Their Territories? When Maps *Territorialize*

At this point, we have enough tools in the toolbox and we have gone through enough discussions of various examples to tackle whether maps "create," in some sense, their territory, as opposed to just representing it.[3] In particular, it is appropriate to address the issue after our discussion of organizations, which indicates the plausibility of map-like items in the plausible absence of the intended territories.

Some may want to consider the issue as a non-starter. The untutored intuitions and practices of map users and mapmakers, after all, converge on the idea of a map-independence of territories: here is *the top of the mountain* and down there is a *village* on the *river*: let us just put there three items on the map, keeping in mind the systematicity constraint.

Maps do work by abstraction: they only represent features of the territory that are of some interest to the mapmaker. They sin by omission but also of commission, if they are mistaken (on purpose, or not). But all

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

of this presupposes a territory has features that can be singled out and abstracted from, a territory that acts as a normative horizon for us to consider the map as accurate or mistaken.[4] There would be no point in *blaming* a map or a mapmaker for their *mistakes* if we did not accept that a territory was there before its representation. Some notion of truth must be available to the map theorizer, short of making the story of map distortion or manipulation untellable. Similarly, given their abstracting power, maps make us see only certain features of the territory, those the mapmaker wants us to pay attention to. This is an epistemological fact about maps, and it has no ontological implications. From the fact that we only see certain features of a territory, it does not follow that the territory has only those features.

A second, perhaps more important, point: an important tradition claims, on the one hand, that maps can be used as instruments of power and control (they are), but at the same time in so doing it says nothing *specific* about maps. Verbal descriptions of a territory can be instruments of power, control, and epistemic injustice via the same mechanisms used by maps: abstraction, omission, selective attention, and invisibilisation. Still, it is clear that maps have a particular power, which appears to come from the way they package information and deliver it to the visual system. The present project is certainly compatible with postcolonialist and deconstructionist critiques of maps. Two goals are of importance here, both aimed at addressing the specificity of map mechanisms: how can we assess maps' relation with truth, the semantic of maps, to account for the various forms of distortion and so on? And why are maps so distinctly effective in sustaining various territorial agendas?

Here are what I consider to be relatively clear cases of maps that do "create" territories.

First, *directive* maps. These are maps that prescribe how a territory should be. Build a skyscraper there. Enlarge the waterfront there. Tear down that old shingle house. Draw a boundary between two newly recognized countries. Huge consequences will follow. But of course, in the present discussion, these situations are a clear Distraction: the interesting claim about maps that create territories concerns *descriptive* maps, not directive maps. Directive maps have a world-to-representation direction of fit, as opposed to the representation-to-world direction of fit (Searle 1983).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Second, consider names for the days of the week. Arguably, there is no territory they describe, as there is no fact of the matter that makes a certain day a Monday as opposed to a Sunday. This being said, if we all agree to using a calendar, then we all agree on there being Sundays, and if we further agree maybe on something special going on on Sundays (or on any other day, according to the culture you pick), then the territory of temporal events *will* change—there will be different driving patterns on Sundays, people will do different things during that day, and so on, and some of these changes will actually get sanctuarized in contracts and legal texts. Calendars are cognitive artifacts that make interpersonal coordination possible and inevitably induce coordination, as they make solutions to coordination games salient, thereby nudging those games.[5]

Third, Borgesian maps territorialize—almost trivially. They change the territory to the point of at times hiding the territorial features they represent or highlight, in some cases even replacing them. Painting street marks changes the visual aspect of the street. This confirms that it pays, perhaps, to consider Borgesian maps as a type of conceptually (not historically) prototypical map, from which other more mundane types are derived. Territorialization becomes a straightforward feature of maps, to the point that we may need to indicate departures from territorialization in non-Borgesian maps.

(Distraction: Territorialization becomes an easy game in the digital age. Electronic dog collars send an electric shock when the dog trespasses a preset frontier.)

But if non-Borgesian maps created their territories, what would the *mechanics* of this creation be like? Consider the following example. The "territorialization" of maps may occur whenever factual information on a map determines a behavior, in response to its use, that ends up modifying the territory. A census map may digitize areas of different but sufficiently similar income, representing them as discrete shades of gray. Looking at this gray-scaled map, people may make choices, such as moving to a different neighborhood, and this in turn ends up hardening the census distribution of income segregation between communities.[6] Where there once was, in the territory, a soft transition that the map rendered in discrete levels, there now is a series of discontinuities in census, with well-defined boundaries. This practice, known as redlining, that designates areas as "undesirable" to certain populations has led to making those areas even less desirable. For instance, once redlined, the areas attract less investment (e.g., in terms of public green) (Hoffman, Shandas, and Pendleton 2020; Nelson, Marciano,

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

and Connolly 2020). The uniform census levels, with abrupt transitions between areas, have become real. The map became territorialized.

But this mechanism is a dynamic evolution made possible by the original descriptive use of the map. Had there been no geographically distributed census differences in the first place, there would have been no point in mapping them. The mapping then hardens the differences, the map territorializes, but the starting point of the process is an independently existing territory.

Maps and Disorientation

At the end of our exploration, I'd like to broach a question that has consumed researchers and the general public in the last decade. Are we losing our wayfinding abilities because of our over-reliance on GPS-assisted devices and the diminished use of paper maps? Are we delegating too much of our mental life to devices that we use in M4 mode? This may well be the case (Ruginski et al. 2019; Dahmani and Bohbot 2020), although arguably the loss is, in the end, compensated by the many improvements in daily life that are made possible by massive geo-localization. I take the latter to be an open question, both on the factual and the normative side; I have no recommendations to make. Other issues, though, seem to have attracted less attention. We can get input from a GPS-assisted device without having to look at a screen (and, when we drive, we must); a synthetic voice imparts directions. This means that although the devices can and do display maps, we end up looking less and less at them. And when we do look at maps, it's a miniature map that we use. Sure enough, we can zoom in and out at will, shrinking the representation of the whole county so that it fits into the few square centimeters of the display. In contrast, paper maps supported a fully different and richer set of operations, such as looking at the detail of the bay and then placing the detail in the larger context of the county, keeping both bay and county at scale, giving us a sense of progress as we move around. But the main difference between a digital and a paper map is that when we use a paper map, we need to find our position on it (which is the position of the map in the territory), and this is only possible if we engage *actively* with *both* the map *and* the territory. This engagement is a form of inquiry that is fueled by interpretation, hypotheses, inferences, and trials with the paper map. We identify landmarks in the territory based on facts

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

we read on the map, and conversely, we find representations of landmarks on the map based on what we observe of the territory. GPS-assisted technologies relieve us from this admittedly complex, at times fastidious, work that some people are able to perfect. We may still look around as we drive in a new landscape, but the *need* to look at the territory is no longer there. And if there is no longer a need, we do not observe the territory actively anymore. And a territory we no longer look at is a place we lose.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Appendix: Maps in the Brain?

Map generation in the brain is a kind of preemptive domestication. Once the brain discovers that you can collect information through maps, other ways of collecting information become hard to discover. But not all information is organized in map-like fashion, not even spatial information.

Retinal Images as Maps

Retinal "images" are best construed as maps; as the retinal sheet is a transducer, it maps light distribution. The retina does not know much about bananas and clouds, but it does a good mechanical job at mapping the light from clouds and bananas. It is, after all, a proximal Albertian Window.

Mental Maps: Shortcuts and Detours

What about the construction of a *mental* map of the environment? These are not directly generated by sensory maps, they are more complex, and we do not even know if their vehicle is map-like. There are several reasons to postulate the existence of mental maps (Tolman 1948; Dudchenko 2010). One is the fact that we can plan *shortcuts* and *detours* through places that are not currently observed. Another is that we can document distortion in mental representations (e.g., placing Venice east of Palermo) that are best explained if we postulate operations on a map-like mental vehicle (Tversky 1981).

Many recording devices (in the brain) also use maps because they are a simple solution to getting spatial information registered spatially. Sheets of captors are a simple (biological) device. They have the double advantage of ease of capture of the spatial information, and of encoding the captured information in a user-friendly way. The neo-cortex is replete with map-like

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

structures of receptors such as retinotopic maps: (Brewer et al. 2005; Rosa 2002) and maps of body parts. The "homunculus map" of Penfield and Boldrey (1937) is probably the most graphic example of a cortical map.

Place Cells Create Mappings, Not Maps

However, in the case of place cells in the hippocampus—the most important of all brain devices that underscore our navigational abilities—there is evidence that they are not structured in a map-like fashion (Kentros 1998).

Discussing the role of place cells in the hippocampus, and quoting John O'Keefe, Bond (2020) writes,

> The place cells in the hippocampus—as opposed to their place fields—certainly look nothing like a map: place cells that are next to each other don't necessarily represent adjacent locations on the ground, and the way the brain assigns place fields to place cells seems quite random. Furthermore, the entire array gets shuffled around—or "remapped," as neuroscientists say—whenever the animal enters a new room. So far, no one has been able to predict how the place cells will behave from one new space to another, or where their place fields might be.
>
> "The lack of a topographical structure in place cells has always been a bit embarrassing to me," says John O'Keefe. "I've worked in an anatomy department my whole career. If you look in the neocortex at the cells that represent a finger, they're sitting alongside the cells that represent the next finger, so there is some topographical representation. But when you have a structure that doesn't do that, where one place cell is sitting far from the place cell that's representing the place next to it, and it's meant to be a map . . . It ain't a map." (45)

So, we learn that the main brain device meant to deliver spatial orientation is not organized like a map. A contrasting account is given in Whittington et al. (2019).

Are There 3D maps in the Brain?

We accepted that there are 3D maps of 3D territories. Does the brain use any? Let me clarify. The same-dimensionality constraint has it that maps are typically 2D vehicles (e.g., pieces of paper, sheets of cells, displays of computers) that represent 2D territories (e.g., surfaces). Of course, they can code 3D features of the territories (e.g., isolines), but their basic work is to register properties of superficial parts of the territory. (Another way to put this is that we should distinguish between being 2D and being flat.) Now imagine having a 3D model of a mountain, such that if you cut through the

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

model you can see in different colors the geological strata and so on. This cake is a 3D map of a 3D territory. Is there anything similar in the brain?

According to neuroscientist Patrick Cavanagh,[1]

> There is no 3D representation that anyone has found where volume is represented. It seems that everything is a surface map and each surface point can have a depth value as a feature, like it might have a color value. That is called a depth map. It seems unlikely that anything other than the visible surface is represented, nothing for the light streaming through space in all direction between the observer and the first visible surface and nothing for the materials underlying that surface.
>
> We can reason about paths through space with obstacles and about the back side of an object (we would be surprised if it weren't there when we went around the back to look) and about the interior of objects we see (jug of milk vs bowling ball). So maybe somewhere there is a volumetric representation of the scene in front of us for us to reason about non-visible aspects, but no one has found that yet. We may do it through other means and we may not do it so well.
>
> The main hunt of the last 30 years was for a map in world coordinates as opposed to retinal coordinates, one that didn't update every time we moved our eyes or body. That too has failed to turn up anything.

So, in the brain there are no 3D maps of 3D territories. One hypothesis is that this peculiar architecture is heavily constrained by the structure of the receptors/transducers. Biological structures find it very convenient to arrange them in sheets, and, well, they do not have many other options. All subsequent information processing and model building goes through the sheet bottleneck, format-wise: sheet in, sheet out. Actually, the sheet organization is very convenient because it gives you a map with very little added costs.

Engineers created other types of receptors. As we saw in chapter 3, some particle detectors are built around piles of sheets of wire in the three dimensions so that when a particle ionizes the gas in the chamber, the nearest wires register the position as in a 3D naval battle, and one can then reconstruct the trajectory in 3D. But animals do not seem to have this. Is there anything that comes close to it? Suppose you hold a marble in your semi-clenched hand. The marble moves around, and we can feel its position in 3D space. Or we can explore the surface of convex objects by covering them with a concave-shaped hand. But these are approximations, and here it is reasonable to suppose that there is a lot of interpolation of 2D-organized data. Two-dimensional internal maps are strongholds of mental space representation, which may partly explain the cognitive appeal of the 2D external maps we keep on using.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Notes

Chapter 1

1. As theoretical posits, internal representations can further be taken literally or reinterpreted as glosses (Egan 2014) that frame empirical research or facilitate cross-disciplinary dialogue. The choice, however, is not directly relevant to our investigation.

2. Note that if some photographs, such as satellite photographs, come close to being taken as maps, this is because the viewpoint is so remote as to render the representation blind to parallax limitations. Hence, the satellite photograph is (almost) portable.

3. Thanks to Pablo Fernández-Velasco for pointing this out.

Chapter 2

1. Not only is it beyond the point to make a comprehensive list of the relevant publications on maps from different disciplinary quarters but it is unpractical to even make a list of the texts I have looked at when writing this book. I single out three book-size background readings that I consider particularly important for the present project: MacEachren (1995), Tyner (2014), and Jacob (1992).

2. This contrasts with the claim made in Greenberg (2021) to the effect that "projection from a viewpoint is the central principle of map semantics, just as it is for pictures."

3. Note that "flat," in this case, does not mean "horizontal."

4. See chapter 5 for a discussion of globes.

5. We shall discuss some in chapter 5.

6. See chapter 7.

7. Janek Guerrini, personal communication, 2022.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

8. As everyone has experienced, Google Maps and like services are best conceived of as visual supports for contextually delivered advertisement. They can hardly be used for navigation without the Google software. In the 2021 version of Google Maps, it is hard to even find street names. Google maps nowadays seem to do the following: (1) locate you on a rough map; (2) advertise the presence of a potentially interesting business in the neighborhood; and (3) let you input a request for getting you from your current location to that potentially interesting business. Although step (2) is spatial, it almost has the form of an existential statement ("there is a potentially interesting item nearby," as the map's roughness (1) makes it difficult to find your way except through (3). We constantly negotiate the absence intuition when using these types of spatial representations.

9. More should be said about the exact nature of metro maps, which are hybrid representations: part map, part diagram. See chapter 9.

10. And, of course, the geo-referencing of ancient maps is a complex endeavor. For more on this, see Dempsey (2011).

11. This point echoes Wittgenstein who states: "In our notations there is indeed something arbitrary, but this is not arbitrary, namely that if we have determined anything arbitrarily, then something else must be the case. (This results from the essence of the notation)" in (*Tractatus* 3.342).

12. More about occlusion in chapter 3.

13. See chapter 3 for a discussion of this point.

14. As Gilbert Ryle once wrote, a person "was ridiculed when he drew a map in which the contour lines were crossed or broken" (Ryle 2009, 282).

15. The term "mental map" has many uses. In human geography, it has been popularized by the work of Gould and White (1992), which however appears to refer to a way to map the preferences of individuals about a certain territory.

Chapter 3

1. Photographs are static pictures, but nothing momentous depends on their static nature. When you open your camera app on your smartphone, you see in the display a moving picture of the scene in front of the camera. This item, the moving picture, is constantly regenerated on the screen and not saved in the camera's memory. It is a volatile picture. No reason to deny it pictorial status.

2. We may want to stabilize the terminology and say that maps map a *territory* and pictures depict a *scene*.

3. See Pirenne (1970) and Benovsky (2010). See Casati (1991) for the discussion of photorealistic pictures as paradigmatic.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

4. Except in the limit case of pictures of 2D items, particularly pictures of maps. See figure 3.3 for the discussion of the limit case of *maps of maps*.

5. See figure 4.4 for a discussion of the use of pictures for navigation.

6. On the logic of occlusion, see Santos et al. (2015)

7. See Palmer (1999). See also Casati (2019b) for some twists of this in the case of line drawings, best construed as maps of the light distribution of occlusion boundaries.

8. For a discussion of the import of the mirror "image" being at the surface of the mirror, see Casati and Cavanagh (2019).

9. We may need a name for the sound equivalent of Albertian Window: the "Helmholtzian Window" perhaps?

10. Some homework: rework the account of pictoriality as *seeing-in* in terms of seeing a map and then seeing a picture in the map one sees.

11. See Chaput (2019). Casati and Cavanagh (2019) call "faux shadows" reflectance distributions that mimic the luminance distribution of shadows.

12. In northern spring and summer, some areas within the northern polar circle will see the midday Sun in the south, and in southern spring and summer, some areas within the southern polar circle will see the midday Sun in the north.

13. On the derivation of maritime maps from verbal descriptions, see Alves Gaspar (2010).

14. Once more, the Wittgensteinian point of *Tractatus* 3.342.

15. See chapter 8 on memorized lists such as the list of the letters in the alphabet.

16. A piece of text with section titles qualifies as a Borgesian map (see chapter 4).

17. Radar and sonar maps (yes, *maps*, not images!) are generated automatically by measuring the time elapsed from the emission to the echo of relevant signals and by then plotting time as distance from the source at the appropriate angle. The production mechanism is hindered by occlusion (occluders block propagation of the signal) and thus the maps display specific shadows (Raynal et al. 2014).

18. I'm indebted to Francesco Pierini for discussions of this point.

19. J. Droulez, personal communication.

20. See chapter 5 for an account of astrolabes as perspectival-computational models.

21. Shin (1994) is the most thorough discussion of the logic of Peirce's diagrams.

22. See Giardino (2013) for an account of the use of diagrams.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Chapter 4

1. Distraction: States may claim underwater territories that stretch far beyond their coastlines, thereby falsifying, at times, the disconnection inference we can draw from maps that do not care about underwater boundaries. Remember to choose the right map for your inferences!

2. As of February 2023, I could not find, among the Unicode symbols, the symbol for the map pin. I proposed its addition to the Unicode consortium. Maybe the writing of this book will have some practical effects.

3. See chapter 3.

4. The recent trend of sensible cartography or environmentally sensitive cartography is neatly exemplified by marine maps.

5. A version of the following sections of this chapter was written with John Kulvicki and John Zeimbekis, and appeared in *Analytic Philosophy*, 2020. I modified the text to insert it in the chapter, here reproduced with the agreement of the authors and of the journal.

Chapter 5

1. A separate discussion is the status of GPS-generated representations. The GPS creates a new physical layer (of perturbations of the electromagnetic field around the Earth) that is a second physical "world." GPS captors only "see" this second world, which is not a model of Earth and has the convenient property of being mathematically well behaved. The tricky part is to coordinate properties of this second world with geographic features of Earth. A GPS representation, in the end, can only be as good as the cartography it uses. Of course, things are improving as the cartography itself is more and more GPS-generated. See Giudici (2016).

2. Presumably, that is, *ceteris paribus*. Information about the main wind patterns can make you change your assessment of travel time.

3. All the globes we buy carry metric information. John Kulvicki (personal communication) suggests that there could be nonmetric globes.

4. I once briefly commented on a *New York Times* article claiming that "distinct from the geographic North Pole, where all the lines of longitude meet at the top of the world, the magnetic pole is the point that a compass recognizes as north [. . .] On multiple occasions during Earth's long history, the magnetic field has weakened dramatically. The north magnetic pole slipped toward the bottom of the planet, and the south magnetic sauntered toward the top. The process took a few thousand years, but by the time the field's full strength returns, it has flipped." The comment: Hmm, is the South of the planet its *bottom*, the North its *top*? For the full article, see Hall (2019).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

5. For the pedagogical advantages of the situated globe, see Casati (2013). See also a rich program for using the globe in primary and secondary schools that has been designed by Nicoletta Lanciano and collaborators (Lanciano 2020).

6. Some historical globe sundials have been documented. See Fantoni (1989).

7. Recall what was said in chapter 2 about a map and its automatic assignment. If in Paris you encounter a faceless globe, then you can reverse engineer a map of the Earth on that globe knowing Paris's latitude, knowing that Paris has the topmost location on the faceless globe, and finding out about the meridian of Paris.

8. Are spatial models taken to be factive the way maps are?

Chapter 6

1. For the historical and cultural aspects of horology, see Hill (1996), Landes (1985), and Cipolla (2003). The account of clocks presented here draws from Casati (1997).

2. (Longo 1999; Smith 1988). The issue may be terminological to a point. A clock's ticking mechanism cannot be simulated by a Turing machine. However, timers and stopwatches can be programmed, so in the end a clock, properly set, can in fact incorporate a computer.

3. What is the input of rod measurement? It is an interval; the placed rod outputs its length if it is inscribed with measurement marks (0/1/2/3/4/5/6/ . . .).

4. The centered spatial web belongs to the family of Borgesian maps.

5. Another interesting hybridization of the two functions was the Parisian countdown timer displayed at the Beaubourg art museum in Paris, which counted the number of seconds separating the then present (last century) from the year 2000. Its function was parasitic on the setting of the counter by using a standard clock.

6. In chapter 7, we discuss in detail raster structures, which clockfaces are an example of.

7. Technically, this only holds for the 24-hour clockface. For an ordinary 12-hour clockface, we would need two stretched half-maps of the world, taking turns. A small design challenge for clockmakers.

8. The Paris meridian, traced on the ground in various locations across France, is Borgesian.

9. The application and its source code are freely available (with CC BY-NC-SA 4.0 creative commons licensing) on online app platforms and code databases. See https://play.google.com/store/apps/details?id=casa.lambda.chronogeoscope&hl=en_US. The source code is available at: https://github.com/webplate/chronogeoscope.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

10. Like the situated globe, the Chronogeoscope is a *slow* educational device. It spins at the angular speed of the Earth (15° per hour), and thus your vision can't tell that it is in motion. You have to wait and use your memory to realize the change. Sure enough, the general educational trend is toward optimizing learning time ("the quicker, the better"). Given the nature of the tool, we cannot afford time shortcuts (though, of course, one can shoot a time-lapse of the web page or of the app to see Earth's rotation in warped time.) Here is a suggestion for educators who wish to use the Chronogeoscope: display it in the morning, when class begins, and then go back to it later in the day. Compare what is displayed with the movement of an ordinary clock's hour hand, spinning twice as fast as the Chronogeoscope's 24-hour dial. Take screenshots and post them to the walls. Allow students to comment on the differences they see. And if you have a situated globe in your courtyard, compare its readings with those of the Chronogeoscope.

Chapter 7

1. One may contrast music notation with statistics notation. Statistics notation was designed by a community of mathematicians, arguably a priori in favor of clarity; it has a short history, less than half a century to stabilize symbols; and it expresses a handful of notions, a microscopic fraction compared to the variety of music referents. And yet, this notation is unfriendly. "P" means three different things in different contexts; sometimes you use a circumflex and sometimes a horizontal bar to say the same thing. Why?

2. "Score" is used here as a synonym of "sheet music."

3. We are abstracting here from the many complex temporal properties of spoken words. Just imagine someone spelling the words slowly so as to articulate the various phonemes.

4. A minimal usable lexicon would use a single type of note to indicate a unit duration, and the tie to bind together notes. As I previously stated, we opt for a system that, although simplified relative to SMN, does not depart too much from SMN. This choice allows us to explore some aspects of SMN. It should be noted that the minimal lexicon would be rather different from SMN, whose class of symbols is open-ended: in principle, there is no shortest describable unit in SMN, as the represented durations are ever-decreasing fractions of the unit note.

Chapter 8

1. Master navigator Mau Piailug can be heard to recite the sequence of names of the stars used by Satawal navigators at about minute 42 in Low (1983).

2. (Gallistel 1993; Karmiloff-Smith 1992).

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

3. (Fodor and LePore 1992).

4. Constraints on the "processability" of S may not be relevant: it *does take time* and effort to learn numerals.

5. (Merzbach and Boyer 2011), although this is only a speculation so far as I can ascertain. Another problem concerns the use of numerals in illiterate cultures; one could look into where these communities got their numerals from, for instance—hypothetically—from a nearby culture that it in turn incorporated it from written language.

6. If bees can dance a map (Morizot 2020), humans can speak out a map.

7. It is possible that spoken language incorporates many other nonlinguistic elements, pieces of the external world that we learned to pronounce. Possibly demonstratives are examples of such incorporation—they are pronounceable gestures.

8. Individual piano keys are numbered by piano manufacturers, representing their position on the keyboard.

9. Here as elsewhere, we could indulge an analysis of many examples to assess their map-like structure. For instance, a roll call list to register attendance is not a map because the class is not an organization; but maybe the class is an organization when its members sit in alphabetical order!

10. A design solution: in waiting rooms, number the chairs! Who comes first occupies chair number 1, and so on. If the number of customers does not exceed the number of chairs, there is always a possibility to occupy chair number 1 again; this works a bit like the marked box that allows you to keep track of medications taken during the week.

11. We suspect that things with queues are strange when we see ways to regulate access to service. Dutch museums experimented with paying the ticket at the exit, not at the entrance. In this way, visitors are not discouraged from entering the museum; and they can decide to leave according to the length of the exit queue, possibly staying in the museum longer and maybe even shopping at the bookstore. Why did it take so long to find this solution? The rules of the queue as an institution seem to provide that if the place in the queue is an indication of merit to get something, it will seem strange that you queue *after* getting what you wanted.

Chapter 9

1. The project of an alternate Venice map was done at University IUAV of Venice with the collaboration of Stéphanie Roisin and Magda Stanová. It evolved into a publication (Casati, Stanová, and Roisin 2017), some of whose passages are reprinted here with permission.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

2. In the controversial practice of gerrymandering, very odd shapes are generated, but it is seldom proposed to create districts with parts that are disconnected (Casati 2001).

3. I give here a marginal contribution to a very large and complex discussion. See Searle (1995) for the general problem, as well as the discussion in MacEachren (1995) and in particular in Winther (2020).

4. An interesting case is that of renaming items in a territory—the most dramatic example of which is the massive colonial renaming of the Americas. See Bond (2020).

5. Calendars are halfway between non-map organizers and map-like organizers. (See the discussion of clocks and clockfaces in chapter 6.)

6. I was made aware of this by David Mark.

Appendix

1. Personal communication, March 13, 2019.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

References

All quoted websites have been consulted in April 2023.

Aguilera, Mariela. 2016. "Cartographic Systems and Non-linguistic Inference." *Philosophical Psychology* 29, no. 3 (January): 349–364. https://doi.org/10.1080/09515089.2015.1089431.

Alves Gaspar, Joaquim. 2010. "From the Portolan Chart of the Mediterranean to the Latitude Chart of the Atlantic: Cartometric Analysis and Modeling." PhD Diss., Universidade Nova de Lisboa.

Ashworth, Mick. 2019. *Why North Is Up: Map Conventions and Where They Came From*. Chicago: University of Chicago Press.

Baldwin Bowsky, M. W., and V. Niniou-Kindeli. 2006. "On the Road Again: A Trajanic Milestone and the Road Connections of Aptera, Crete." *Hesperia* 75, no. 3 (July–September): 405–433. https://doi.org/10.2972/hesp.75.3.405.

Baum, L. Frank, and Barry Moser. 1899/1986. *L. Frank Baum's The Wonderful Wizard of Oz*. Berkeley: University of California Press.

Benovsky, Jiri. 2010. *Qu'est-ce qu'une photographie*. Paris: Vrin.

Bloom, Paul. 1998. "Theories of Artifact Categorization." *Cognition* 66, no. 1 (April): 87–93. https://doi.org/10.1016/S0010-0277(98)00003-1.

Blumson, Ben. 2014. *Resemblance and Representation: An Essay in the Philosophy of Pictures*. Open Book Publishers.

Bond, Michael. 2020. *Wayfinding: The Art and Science of How We Find and Lose Our Way*. Picador.

Borges, Jorge Luis. 1944/1998. *Collected Fictions*. New York: Viking.

Bray, Hiawatha. 2014. *You Are Here*. New York: Basic Books.

Brewer, Alyssa A., Junjie Liu, Alex R. Wade, and Brian A. Wandell. 2005. "Visual Field Maps and Stimulus Selectivity in Human Ventral Occipital Cortex." *Nature Neuroscience* 8, no. 8 (July): 1102–1109. https://doi.org/10.1038/nn1507.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Bronner, Ben. 2015. "Maps and Absent Symbols." *Australasian Journal of Philosophy* 93, no. 1 (August): 43–59. https://doi.org/10.1080/00048402.2014.948463.

Camp, Elisabeth. 2007. "Thinking with Maps." *Philosophical Perspectives* 21, no. 1 (December): 145–182. https://doi.org/10.1111/j.1520-8583.2007.00124.x.

Camp, Elisabeth. 2018. "Why Maps Are Not Propositional." In *Non-Propositional Intentionality*, edited by Alex Grzankowski and Michelle Montague, Vol. 1., 19–45. New York: Oxford University Press. https://doi.org/10.1093/oso/9780198732570.003.0002.

Carr, Nicholas. 2014. *The Glass Cage*. New York: W.W. Norton & Company

Carrara, Massimiliano, and Daria Mingardo. 2013. "Artifact Categorization. Trends and Problems." *Review of Philosophy and Psychology* 4, no. 3 (June): 351–373. https://doi.org/10.1007/s13164-013-0151-6.

Carroll, Lewis. 1894/1988. *Sylvie and Bruno*. New York: Dover.

Casati, Roberto. 1991. *L'immagine: introduzione ai problemi filosofici della rappresentazione*. Firenze: Nuova Italia.

Casati, Roberto. 1997. "Qu'est-ce qu'une horloge?" *Cahiers de philosophie analytique* 1: 189–219.

Casati, Roberto. 2001. "Cognitive Aspects of Gerrymandering." *Topoi* 20 (December): 203–212. https://doi.org/10.1023/A:1017917109736.

Casati, Roberto. 2003a. *The Shadow Club*. New York: Knopf.

Casati, Roberto. 2003b. "Representational Advantages." *Proceedings of the Aristotelian Society* 103, no. 1 (February): 281–298. https://hal.science/ijn_00000331.

Casati, Roberto. 2010. "Hallucinatory Pictures." *Acta Analytica* 25, no. 3 (January): 365–368. https://doi.org/10.1007/s12136-009-0074-0.

Casati, Roberto. 2011. *Prima Lezione Di Filosofia* [Philosophy as Conceptual Negotiation]. Roma: Laterza. https://hal.science/hal-03937758.

Casati, Roberto. 2012. "Mirrors, Illusions and Epistemic Innocence." In *Perceptual Illusions*, edited by C. Calabi and P. Spinicci, 192–201. London: Palgrave Macmillan UK.

Casati, Roberto. 2013. *Dov'è il sole di notte?* Milano: Raffaello Cortina Editore.

Casati, Roberto. 2014. "Numerals and Word Sequences." In *Mind, Values, and Metaphysics*, edited by Anne Reboul, 327–338. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-05146-8_22.

Casati, Roberto. 2016. "The Structure of Standard Music Notation." In *The Theory and Practice of Ontology*, edited by Leo Zaibert, 187–201. London: Palgrave Macmillan UK. https://doi.org/10.1057/978-1-137-55278-5_10.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Casati, Roberto. 2017. "Review of Resemblance and Representation: An Essay in the Philosophy of Pictures, by Ben Blumson." *Australasian Journal of Philosophy* 95, no. 3 (2017): 603–606. https://doi.org/10.1080/00048402.2016.1207691.

Casati, Roberto. 2019a. "Navigation as Instrumental Negotiation. Report and Observations on a Transatlantic Field Trip and Training on Sailboat *Albatros* (2016–2017)." Italian Academy Preprints, 42. https://italianacademy.columbia.edu/paper/navigation-instrumental-negotiation.

Casati, Roberto. 2019b. "Drawing Problems: Thought in Action." *Adaptive Behavior* 29, no. 5 (July): 441–448. https://doi.org/10.1177/1059712319861551.

Casati, Roberto, and Cavanagh, Patrick. 2019. *The Visual World of Shadows*. Cambridge, MA: MIT Press.

Casati, Roberto, and Jérôme Dokic. 1994. *La Philosophie du son*. Nîmes: Editions J. Chambon.

Casati, Roberto, and Valeria Giardino. 2013. "Public Representation and Indeterminicies of Perspectival Content." In *Enacting Images*, edited by Zsuzsanna Kondor, 111–126. Herbert von Halem Verlag.

Casati, Roberto, and Gianfranco Soldati. 1995. "On the Perception of Abstract Objects." In *Studies in Reference and Meaning*, edited by Karlovy Vary, 89–113. Czech Academy of the Sciences.

Casati, Roberto, Magda Stanová, and Stéphanie Roisin. 2017. "The Pedestrian's Venice." *Magazine X Special Issue*.

Casati, Roberto, and Christine Tappolet. 1987. "Review of F. Schier, Deeper into Pictures." *Studia Philosophica* 48: 220–223.

Casati, Roberto, and Achille C. Varzi. 1999. *Parts and Places: The Structures of Spatial Representation*. Cambridge, MA: MIT Press.

Casey, Edward S. 2002. *Representing Place: Landscape Painting and Maps*. Minneapolis: The University of Minnesota Press.

Cavanagh, Patrick. 1999. "Pictorial Art and Vision." In *The MIT Encyclopedia of the Cognitive Sciences*, edited by R. A. Wilson and F. C. Keil, 648–650. Cambridge, MA: MIT Press.

Cavanagh, Patrick. 2011. "Visual Cognition." *Vision Research* 51, no. 13 (July): 1538–1551. https://doi.org/10.1016/j.visres.2011.01.015.

Chaput, Apolline. 2019. *La Représentation Scientifique Des Nanoimages*. Mémoire de Master. Paris: PSL.

Chiabrando, Filiberto, Vincenzo Di Pietra, Andrea Lingua, Youngsu Cho, and Juil Jeon. 2017. "An Original Application of Image Recognition Based Location in

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Complex Indoor Environments." *ISPRS International Journal of Geo-Information* 6, no. 2 (February): 56. https://doi.org/10.3390/ijgi6020056.

Cipolla, Carlo M. 2003. *Clocks and Culture, 1300–1700*. New York: Norton.

Cotnoir, A. J., and Achille C. Varzi. 2021. *Mereology*. Oxford University Press.

Dahmani, Louisa, and Véronique D. Bohbot. 2020. "Habitual Use of GPS Negatively Impacts Spatial Memory during Self-Guided Navigation." *Scientific Reports* 10 (1). https://doi.org/10.1038/s41598-020-62877-0.

Dehaene, Stanislas. 2009. *Reading in the Brain: The Science and Evolution of a Human Invention*. New York: Viking.

Dehaene, Stanislas, and Laurent Cohen. 2011. "The Unique Role of the Visual Word Form Area in Reading." *Trends in Cognitive Sciences* 15, no. 6 (May): 254–262. https://doi.org/10.1016/j.tics.2011.04.003.

Dempsey, Caitlin. 2011. Historical Geography and GPS. https://www.gislounge.com/historical-geography-and-gis/.

Denis, Michel. 2018. *Space and Spatial Cognition: A Multidisciplinary Perspective*. New York: Routledge.

Di Francesco, Michele. 2007. "Extended Cognition and the Unity of Mind. Why We Are Not 'Spread into the World.'" In *Cartographies of the Mind: Philosophy and Psychology in Intersection*, edited by Massimo Marraffa, Mario De Caro, and Francesco Ferretti, 211–225. Dordrecht: Springer Netherlands. https://doi.org/10.1007/1-4020-5444-0_16.

Dretske, Fred I. 1986. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Dretske, Fred I. 1988. *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.

Dudchenko, Paul A. 2010. *Why People Get Lost: The Psychology and Neuroscience of Spatial Cognition*. Oxford: Oxford University Press.

Egan, F. 2014. How to think about mental content. *Philosoophical Studies* 170, 115–135. https://doi.org/10.1007/s11098-013-0172-0.

Efrat-Treister, Dorit, Arik Cheshin, Dana Harari, Shira Agasi, Hadar Moriah, Hanna Admi, and Anat Rafaeli. 2019. "Correction: How Psychology Might Alleviate Violence in Queues: Perceived Future Wait and Perceived Load Moderate Violence against Service Providers." *PLOS ONE* 14, no. 7 (July): e0220395. https://doi.org/10.1371/journal.pone.0220395.

Evans, Jonathan St. B. T. 2003. "In Two Minds: Dual-Process Accounts of Reasoning." *Trends in Cognitive Sciences* 7, no. 10 (October): 454–459. https://doi.org/10.1016/j.tics.2003.08.012.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Evans, Jonathan St. B. T. 2012. "Dual-Process Theories of Deductive Reasoning: Facts and Fallacies." In *The Oxford Handbook of Thinking and Reasoning*, edited by Keith J. Holyoak and Robert G. Morrison, 115–133. New York: Oxford University Press.

Fainstein, Susan S., and Scott Campbell, eds. 2016. *Readings in Planning Theory*. Chichester: John Wiley.

Fantoni, G. 1989. "Due orologi solari greci: i globi di Prosymna e di Matelica." *Rivista Di Archeologia* Supplementi (9).

Fernández Velasco, Pablo, and Roberto Casati. 2020. "Subjective Disorientation as a Metacognitive Feeling." *Spatial Cognition and Computation* 20, no. 4 (May): 281–305. https://doi.org/10.1080/13875868.2020.1768395.

Flammarion, Camille. 1880. *Astronomie populaire: Description générale du ciel*. Paris: Marpon et Flammarion.

Fodor, Jerry A. 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.

Fodor, Jerry A. 2009. Where Is My Mind? *London Review of Books*, February 12, 2009. https://www.lrb.co.uk/the-paper/v31/n03/jerry-fodor/where-is-my-mind.

Fodor, Jerry A. 2008. *LOT 2: The Language of Thought Revisited*. Oxford, New York: Oxford University Press.

Fodor, Jerry A., and Ernest LePore. 1992. *Holism: A Shopper's Guide*. Oxford: Blackwell.

Galilei, Galileo. 2013. *The Starry Messenger*. Delray Beach, FL: Levenger Press in association with the Library of Congress.

Gallistel, Charles R. 1993. *The Organization of Learning*. Cambridge, MA: MIT Press.

Giardino, Valeria. 2013. "A Practice-Based Approach to Diagrams." In *Visual Reasoning with Diagrams*, edited by Amirouche Moktefi and Sun-Joo Shin, 135–351. Basel: Springer Basel. https://doi.org/10.1007/978-3-0348-0600-8_8.

Gibson, James J. 1978. "The Ecological Approach to the Visual Perception of Pictures." Leonardo 11, no. 3 (Summer): 227–235. https://doi.org/10.2307/1574154.

Giudici, Sergio. 2016. *Fare il punto Una storia a ritroso della localizzazione dal GPS a Tolomeo*. Milano : Mondadori università.

Gladwin, Thomas. 1974. *East Is a Big Bird: Navigation and Logic on Puluwat Atoll*. Cambridge, MA: Harvard University Press.

Gombrich, Ernst H. 1960. *Art and Illusion*. London: Phaidon.

Goodman, Nelson. 1968. *Languages of Art: An Approach to a Theory of Symbols*. Indianapolis: Hackett.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Goody, Jack. 1978. *The Domestication of the Savage Mind*. Cambridge: Cambridge University Press.

Gould, Peter, and Rodney White. 1992. *Mental Maps*. 2nd ed. London, New York: Routledge.

Greenberg, Gabriel. 2021. Map Semantics and the Geography of Meaning. http://gjgreenberg.bol.ucla.edu/docs/maps.pdf.

Hagen, Margaret A. 1986. Varieties of Realism: Geometries of Representational Art. Cambridge: Cambridge University Press.

Hall, Shannon. 2019. The North Magnetic Pole's Mysterious Journey Across the Arctic. *The New York Times*. Feb. 4, 2019. https://www.nytimes.com/2019/02/04/science/north-magnetic-pole-model.html.

Harley, J. B., and J. H. Andrews. 2002. *The New Nature of Maps: Essays in the History of Cartography*. Edited by Paul Laxton. Baltimore, MD: Johns Hopkins University Press.

Harrison, Jim. 1994. *Legends of the Fall*. New York: Delta/Seymour Lawrence.

Haugeland, John. 1981. "Analog and Analog." *Philosophical Topics* 12 (1): 213–225. https://doi.org/10.5840/philtopics198112150.

Haugeland, John. 1991. "Representational Genera." In *Philosophy and Connectionist Theory*, edited by W. Ramsey, S. P. Stich, and D. E. Rumelhart, 61–90. Psychology Press.

Haugeland, John. 1998. *Having Thought: Essays in the Metaphysics of Mind*. Cambridge, MA: Harvard University Press.

Healey, Christopher, and James Enns. 2012. "Attention and Visual Memory." *IEEE Transactions on Visualization and Computer Graphics* 18, no. 7 (July): 1170–1188. https://doi.org/10.1109/TVCG.2011.127.

Hecht, Heiko, Marco Bertamini, and Matthias Gamer. 2005. "Naive Optics: Acting on Mirror Reflections." *Journal of Experimental Psychology: Human Perception and Performance* 31, no. 5 (October): 1023–1038. https://doi.org/10.1037/0096-1523.31.5.1023.

Heersmink, Richard. 2013. "A Taxonomy of Cognitive Artifacts: Function, Information, and Categories." *Review of Philosophy and Psychology* 4, no. 3 (May): 465–481. https://doi.org/10.1007/s13164-013-0148-1.

Heim, Arnold. 1913. *Monographie der Churfirsten-Mattstock-Gruppe*. Zweiter Teil, Stratigraphieder Mittleren Kreide. Bern: A. Francke.

Hill, Donald. 1996. *A History of Engineering in Classical and Medieval Times*. London, New York: Routledge.

Hoffman, Jeremy S., Vivek Shandas, and Nicholas Pendleton. 2020. "The Effects of Historical Housing Policies on Resident Exposure to Intra-Urban Heat: A Study of 108 US Urban Areas." *Climate* 8, no. 1 (January): 12. https://doi.org/10.3390/cli8010012.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Hopkins, Robert. 1995. "Explaining Depiction." *The Philosophical Review* 104, no. 3 (July): 425–455. https://doi.org/10.2307/2185635.

Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.

Hutchins, Edwin. 1999. "Cognitive Artefacts." In *The MIT Encyclopedia of the Cognitive Sciences*, edited by R. A. Wilson and F. C. Keil, 126–128. Cambridge, MA: MIT Press.

Huth, John Edward. 2013. *The Lost Art of Finding Our Way*. Cambridge, MA: Harvard University Press.

Hyman, John. 2006. *The Objective Eye: Color, Form, and Reality in the Theory of Art*. Chicago: University of Chicago Press.

Jacob, Christian. 1992. *L'empire des cartes: Approche théorique de la cartographie à travers l'histoire*. Paris: A. Michel.

Jarrett, Keith. 1991. *The Köln Concert*. Mainz: Schott.

Kahneman, Daniel. 2013. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Karmiloff-Smith, Annette. 1992. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.

Kennedy, J. M. 1975. "Drawing Was Discovered, Not Invented." *New Scientist* 67: 523–527.

Kentros, C. 1998. "Abolition of Long-Term Stability of New Hippocampal Place Cell Maps by NMDA Receptor Blockade." *Science* 280, no. 5372 (June): 2121–2126. https://doi.org/10.1126/science.280.5372.2121.

Kirk, Ulrich, Martin Skov, Mark Schram Christensen, and Niels Nygaard. 2009. "Brain Correlates of Aesthetic Expertise: A Parametric FMRI Study." *Brain and Cognition* 69, no. 2 (March): 306–315. https://doi.org/10.1016/j.bandc.2008.08.004.

Kleffner, Dorothy A., and Vilayanur S. Ramachandran. 1992. "On the Perception of Shape from Shading." *Attention, Perception, and Psychophysics* 52, no. 1 (July): 18–36.

Kulvicki, John. 2006. *On Images: Their Structure and Content*. Oxford, New York: Clarendon.

Kulvicki, John. 2010. "Knowing with Images: Medium and Message." *Philosophy of Science* 77, no. 2 (2010): 295–313. https://doi.org/10.1086/651321.

Kulvicki, John. 2014. *Images*. London, New York: Routledge.

Kulvicki, John. 2015. Maps, Pictures, and Predication. *Ergo* 2 (7). https://doi.org/10.3998/ergo.12405314.0002.007.

Kulvicki, John. 2020. *Modeling the Meanings of Pictures*. New York: Oxford University Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Kyselka, Will. 1987. *An Ocean in Mind*. A Kolowalu Book. Honolulu: University of Hawaii Press.

Lanciano, Nicoletta. 2020. *Strumenti per i giardini del cielo: Materiali per le classi, per i musei, per i parchi, per la formazione degli insegnanti e degli animatori culturali*. Bergamo: Edizioni Junior.

Landes, David S. 1985. *Revolution in Time: Clocks and the Making of the Modern World*. Cambridge, London: Belknap Press of Harvard University Press.

Larkin, Jill H., and Herbert A. Simon. 1987. "Why a Diagram Is (Sometimes) Worth Ten Thousand Words." *Cognitive Science* 11, no. 1 (January–March): 65–100.

Legrenzi, Paolo, and C. A. Umiltà. 2011. *Neuromania: On the Limits of Brain Science*. Oxford, New York: Oxford University Press.

Lemaitre, Guillaume, Olivier Houix, Frédéric Voisin, Nicolas Misdariis, and Patrick Susini. 2016. "Vocal Imitations of Non-Vocal Sounds." *PLOS ONE* 11, no. 12 (December): e0168167. https://doi.org/10.1371/journal.pone.0168167.

Lin, Allen Yilun, Kate Kuehl, Johannes Schöning, and Brent Hecht. 2017. Understanding "Death by GPS": A Systematic Study of Catastrophic Incidents Associated with Personal Navigation Technologies. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. New York: Association for Computing Machinery, 1154–1166. https://doi.org/10.1145/3025453.3025737.

Lewis, David K. 1971. "Analog and Digital." *Noûs* 5, no. 3 (September): 321–327. https://doi.org/10.2307/2214671.

Lewis, David. 1972. *We, the Navigators: The Ancient Art of Landfinding in the Pacific*. Wellington: Reed.

Longo, Giuseppe. 1999. "The Difference between Clocks and Turing Machines." In *Functional Models of Cognition*, edited by Arturo Carsetti, 211–232. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-015-9620-6_14.

Low, Sam, dir. 1983. "The Navigators: Pathfinders of The Pacific." *Documentary Educational Resources (DER)*. Watertown, MA.

MacEachren, Alan M. 1995. *How Maps Work: Representation, Visualization, and Design*. New York: Guilford Press.

MacFarlane, John. 2015. "Dretske's Theory of Representation." MS.

Madl, Tamas, Ke Chen, Daniela Montaldi, and Robert Trappl. 2015. "Computational Cognitive Models of Spatial Memory in Navigation Space: A Review." *Neural Networks* 65 (May): 18–43. https://doi.org/10.1016/j.neunet.2015.01.002.

Mamassian, Pascal, and Ross Goutcher. 2001. "Prior Knowledge on the Illumination Position." *Cognition* 81, no. 1 (August): B1–B9. https://doi.org/10.1016/S0010-0277(01)00116-0.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Marconi, Diego. 2005. Contro La Mente Estesa. *Sistemi Intelligenti* 17: 389–398.

Marconi, Diego. 2013. "Pencils Have a Point: Against General Externalism about Artifactual Words." *Review of Philosophy and Psychology* 4, no. 3 (May): 497–513.

Marine Casualties Investigative Body. 2013. *Cruise Ship Costa Concordia Marine Casualty on January 13, 2012 Report on the Safety Technical Investigation*. Rome: Ministry of Infrastructures and Transports.

Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.

Merzbach, Uta. C., and Carl B. Boyer. 2011. *A history of mathematics* (3rd ed). New York : John Wiley.

Monmonier, Mark S. 1996. *How to Lie with Maps*. 2nd ed. Chicago: University of Chicago Press.

Morizot, Baptiste. 2020. *Manières d'être vivant: Enquêtes sur la vie à travers nous*. Arles: Actes Sud.

Moyes, Holley. 2002. The Use of GIS in the Spatial Analysis of an Archaeological Cave Site. *Journal of Cave and Karst Studies* 64, no. 1 (April): 9–16. https://legacy.caves.org/pub/journal/PDF/V64/v64n1-Moyes.pdf.

Nelson, R. K., L. Winling, R. Marciano, and N. Connolly. 2020. "Mapping Inequality." *American Panorama*. https://dsl.richmond.edu/panorama/redlining/#loc=5/39.1/-94.58.

Nicolas, David. 2018. "The Logic of Mass Expressions." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2018/entries/logic-massexpress/.

Noë, Alva. 2006. *Action in Perception*. Cambridge, MA: MIT Press.

Oxenbould, Chris, Stan Honey, and Chuck Hawley. 2015. *Volvo Ocean Race Independent Report into the Stranding of Vestas Wind*. Volvo Ocean Race Organization.

Pacherie, Elisabeth. 1995. "Do We See with Microscopes?" *The Monist* 78, no. 2 (April): 171–188. https://www.jstor.org/stable/27903428.

Palmer, Stephen E. 1999. *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.

Pelletier, F. Jeffry. 1975. Non-singular Reference: Some Preliminaries. *Philosophia* 5 (October): 451–465. https://doi.org/10.1007/BF02379268.

Penfield, Wilder, and Edwin Boldrey. 1937. "Somatic Motor and Sensory Representation in the Cerebral Cortex of Man as Studied by Electrical Stimulation." *Brain* 60, no. 4 (December): 389–443. https://doi.org/10.1093/brain/60.4.389.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Pignocchi, Alessandro. 2010. "How the Intentions of the Draftsman Shape Perception of a Drawing." *Consciousness and Cognition* 19, no. 4 (December): 887–898. https://doi.org/10.1016/j.concog.2010.04.009.

Pignocchi, Alessandro. 2012. *L'oeuvre d'art et Ses Intentions*. Paris: Odile Jacob.

Pirenne, M. H. 1970. *Optics, Painting and Photography*. London: Cambridge University Press.

Plato. 2006. *Phaedrus*. Teddington: The Echo Library.

Pliny. 1967. *The Natural History of Pliny*. London: H.G. Bohn.

Pratt, Ian. 1993. "Map Semantics." In *Spatial Information Theory: A Theoretical Basis for GIS*, edited by Andrew U. Frank and Irene Campari, 77–91. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer. https://doi.org/10.1007/3-540-57207-4_7.

Rankin, Susan. 2018. *Writing Sounds in Carolingian Europe: The Invention of Musical Notation*. 1st ed. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781108368605.

Rankin, William. 2018. *After the Map: Cartography, Navigation, and the Transformation of Territory in the Twentieth Century*. Paperback edition. Chicago: University of Chicago Press.

Raynal, Ann Marie, Douglas. L Bickel, and Armin W. Doerry. 2014. "Stationary and Moving Target Shadow Characteristics in Synthetic Aperture Radar." *Proceedings of SPIE, 9077, n° 90771B-1*.

Rescorla, Michael. 2009. "Predication and Cartographic Representation." *Synthese* 169, no. 1 (May): 175–200. https://doi.org/10.1007/s11229-008-9343-5.

Rey, H. A. 2006. *The Stars: A New Way to See Them*. Boston: Houghton Mifflin.

Robinson, Arthur Howard, and Barbara Bartz Petchenik. 1976. *The Nature of Maps: Essays toward Understanding Maps and Mapping*. Chicago: University of Chicago Press.

Rosa, M. G. P. 2002. "Visual Maps in the Adult Primate Cerebral Cortex: Some Implications for Brain Development and Evolution." *Brazilian Journal of Medical and Biological Research* 35, no. 12 (December): 1485–1498. https://doi.org/10.1590/S0100-879X2002001200008.

Roskies, Adina L. 2007. "Are Neuroimages like Photographs of the Brain?" *Philosophy of Science* 74, no. 5 (December): 860–872. https://doi.org/10.1086/525627.

Royce, Josiah. 2012. *World and the Individual*, Vol. 1. Nabu Press.

Ruginski, Ian Tanner, Sarah Creem-Regehr, Jeanine K. Stefanucci, and Elizabeth Cashdan. 2019. "GPS Use Negatively Affects Environmental Learning through Spatial Transformation Abilities." *Journal of Environmental Psychology* 64 (August): 12–20. https://doi.org/10.1016/j.jenvp.2019.05.001.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Ryle, Gilbert. 2009. *The Concept of Mind*. Routledge.

Salway, Benet 2001. "Travel, Itineraria, and Tabellaria." In *Travel and Geography in the Roman Empire*, edited by C. Adams and R. M. Laurence, 22–66. London, New York: Routledge.

Salway, Benet. 2005. "The Nature and Genesis of the Peutinger Map." *Imago Mundi* 57, no. 2 (2005): 119–135. https://doi.org/10.1080/03085690500094867.

Santos, Paulo E., Roberto Casati, Hannah Dee, Carl Schultz, and Mehul Bhatt. 2016. "Eclipse in Occlusion A perspectival mereotopological representation of celestial eclipses." *Proceedings of the 29th International Workshop on Qualitative Reasoning*. https://ivi.fnwi.uva.nl/tcs/QRgroup/qr16/pdf/13Santos.pdf.

Schier, Flint. 1986. *Deeper into Pictures: An Essay on Pictorial Representation*. New York: Cambridge University Press.

Schumann, Robert. 1983. *On Music and Musicians*. Edited by Konrad Wolff and Paul Rosenfeld. Berkeley: University of California Press.

Schwartz, Barry. 1975. *Queuing and Waiting: Studies in the Social Organization of Access and Delay*. Chicago: University of Chicago Press.

Searle, John R. 1983. *Intentionality, an Essay in the Philosophy of Mind*. New York: Cambridge University Press.

Searle, John R. 1995. *The Construction of Social Reality*. New York: Free Press.

Searle, John R. 2018. "The Map and the Territory." In *The Map and the Territory*, edited by Shyam Wuppuluri and Francisco Antonio Doria, 71–78. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-72478-2_4.

Shin, Sun-joo. 1994. "Peirce and the Logical Status of Diagrams." *History and Philosophy of Logic* 15, no. 1 (March): 45–68. https://doi.org/10.1080/01445349408837224.

Smith, Brian Cantwell. 1988. "The Semantics of Clocks." In *Aspects of Artificial Intelligence*, edited by James H. Fetzer, 3–31. Studies in Cognitive Systems. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-009-2699-8_1.

Sorensen, Roy. 2013. "The Twin Towers Riddle." *Philosophical Studies* 162, no. 1 (July): 109–117. https://doi.org/10.1007/s11098-012-9992-6.

Spelke, Elizabeth S., and Sanna Tsivkin. 2001. "Initial Knowledge and Conceptual Change: Space and Number." In *Language Acquisition and Conceptual Development*, edited by Melissa Bowerman and Stephen C. Levinson, 70–98. Cambridge: Cambridge University Press.

Sperber, Dan, and Lawrence A. Hirschfeld. 2004. "The Cognitive Foundations of Cultural Stability and Diversity." *Trends in Cognitive Sciences* 8, no. 1 (January): 40–46. https://doi.org/10.1016/j.tics.2003.11.002.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Sprevak, Mark. 2019. "Extended Cognition." In *Routledge Encyclopedia of Philosophy*, edited by Edward Craig. Routledge. https://doi.org/10.4324/9780415249126-V049-1.

Stanovich, Keith E. 2009. "Distinguishing the Reflective, Algorithmic, and Autonomous Minds: Is It Time for a Tri-Process Theory?" In *In Two Minds: Dual Processes and Beyond*, edited by J. S. B. T. Evans and K. Frankish, 55–88. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199230167.003.0003.

Strawson, P. F. 1959. *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.

Strayer, Hope. 2013. "From Neumes to Notes: The Evolution of Music Notation." *Musical Offerings* 4 (1): 1–14. https://doi.org/10.15385/jmo.2013.4.1.1.

Tabarly, Eric. 1974. *De Pen Duick en Pen Duick: De la transatlantique à la transpacifique*. Paris: Arthaud.

Tolman, Edward C. 1948. "Cognitive Maps in Rats and Men." *Psychological Review* 55 (4): 189–208. https://doi.org/10.1037/h0061626.

Tucker, Holly. 2018. *City of Light, City of Poison: Murder, Magic, and the First Police Chief of Paris*. New York: Norton.

Tüscher, Anne. 2006. "Images en perspectives: La valeur cognitive de l'axonométrie." These de doctorat. Paris: EHESS. http://theses.fr/2006EHES0067.

Tversky, Barbara. 1981. "Distortions in Memory for Maps." *Cognitive Psychology* 13, no. 3 (July): 407–433. https://doi.org/10.1016/0010-0285(81)90016-5.

Tversky, Barbara, M. Agrawala, J. Heiser, P. Hanrahan, and D. Phan. 2007. "Cognitive Design Principles for Generating Visualizations." In *Applied Spatial Cognition*, edited by Gary L. Allen, 53–73. Mahwah, NJ: Lawrence Erlbaum Associates.

Tversky, Barbara, and Paul U. Lee. 1998. "How Space Structures Language." In *Spatial Cognition*, edited by Christian Freksa, Christopher Habel, and Karl F. Wender, 157–175. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-69342-4_8.

Tyner, Judith. 2014. *Principles of Map Design*. New York: Guilford.

Varzi, A., ed. 2001. "The Philosophy of Geography." *Topoi* 20 (2).

Voltolini, Alberto. 2015. *A Syncretistic Theory of Depiction*. New York: Palgrave Macmillan.

Walton, Kendall L. 1990. *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Cambridge, MA: Harvard University Press.

Weisberg, Michael. 2012. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Whittington, James C. R., Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy E. J. Behrens. 2019. "The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalisation in the Hippocampal Formation." *Cell* 183, no. 5 (November): 1249–1263. https://doi.org/10.1016/j.cell.2020.10.024.

Willats, John. 1997. *Art and Representation: New Principles in the Analysis of Pictures*. Princeton, NJ: Princeton University Press.

Winkler-Rhoades, Nathan, Susan C. Carey, and Elizabeth S. Spelke. 2013. "Two-Year-Old Children Interpret Abstract, Purely Geometric Maps." *Developmental Science* 16, no. 3 (May): 365–376. https://doi.org/10.1111/desc.12038.

Winther, Rasmus G.. 2020. *When Maps Become the World*. Chicago: University of Chicago Press.

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Oxford: Blackwell Publishing.

Wood, Denis, and John Fels. 1992. *The Power of Maps*. New York: Guilford Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Index

(Not a map, but a *mapping* of the text! For a *map* of the text, see the table of contents on page v.)

Absence intuition, 4, 52–53
Abstract entities, 87, 187
Abstraction, 3–4, 6, 29–30, 49, 64–66, 99, 115
  abstract use of a map, 120–121, 147
  from viewpoint, 104
Affordances, 20, 29
Affordance perception, 29
Ahead-up, 8
AIS (Automatic Identification System), 99
Albertian Window, 5, 83–85, 87, 91, 94, 99, 105, 130–131
Alpine Sequence, 181
Artifacts
  categorization of, 96
  epistemic, 7
  representational, 61
Artifact words, 28
Aspects, 5–6, 89–90
Assignment, 53, 58, 61
  automatic, 4, 61, 65–66, 79
  structured, 78
Astrolabes, 147–148
  are not maps of stars, 147
  as computational models, 147
ATMs, 26
Atomic expressions, music 178–179
Attention, 13
Augmented reality, 127
Axonometry, 100–105
  long-distance comparison in, 105
Azimuthal equidistant projection, 45–46, 157–163

Bar Limit Principle, 178
Barriers, 35–36
Bars, 177–179
Baum, L. Frank, 122
Biases, 12
  cognitive, 35
  perceptual, 91–94
Bird's eye view, 120
Blocks, 202–208
  can be circumpaced, 202
  (*see also* Islands)
Body, extended, 21–22
Borgesian maps, 6, 79, 121–132, 189, 200, 210
  are pre-oriented, 128
  definition, 122
  no Red Dot needed on, 128–129
  spatially mimetic, 124
  territorialize, 210
  useful, 124–127
Boundaries, 75

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Brain maps, 98
Brain scans, 98–99
Brakes, Cognitive, 35
Bridging, representational, 37–39, 184

Calendars, 151, 210
  as abstract temporal maps, 151
  tear-off, 151, 165
Camouflage, 36
Camp, Elisabeth, 67–68, 71
Cartographic Age, 7
Cartographic software, 83
Cavanagh, Patrick, 215
Ceiling, representational, 38–39, 120–121
Celestial maps, 147, 164–165
  not maps of stars, 105
Chronogeoscope, 161–166
Clock hand, as a meridian on a map, 163
Clockfaces, 1, 7, 81, 151–166
  as maps, 156–165
  as spatial maps, 160
  as visible part of dynamic calendar, 151
Clocks, 7, 151–166
  cannot be simulated by a computer, 152
  Spiraling Clock, 157–158
  Wheel clock, 157–158
Cognition, distributed, 8
Cognition, extended 1
Cognitive advantage, 6, 14–15, 24, 33, 36, 139, 215
  for Southerners, 91–94
Cognitive architecture, 8
Cognitive artifacts, 1–2, 19, 21, 28–30
  classification of, 28–29
  design constraints on, 39–40
  nonrepresentational, 28–29
  representational noncognitive, 28–29
  representational, 28–29
Cognitive engagement, 1, 7
Cognitive loss, 8
Cognitive mechanics, 28–30
Cognitive toolbox, 1 ff.
Color patches on maps, 69
Colors, can overlap on formal maps, 74
Compositionality, 51–52
Computations, 11 ff., 100, 117–118
  core navigational, 18
  core, 23, 55
  navigational, 16
Computational (non)equivalence, 33, 68
Computational tasks, auxiliary, 16
Computational tasks, core, 16
Computational-representational framework, 2 ff.
Concatenations, 176
Concavity/convexity reversal, 92–93
Constellations, 164
  as naturefacts, 163
Content, 29
  of pictures, dual theories of, 90
  bare-bones, 90
  fleshed-out, 90
Coordinate grids, 55, 173
Coordination games, 210
Costa Concordia accident, 8
Costs, cognitive, 182, 188
Costs, inferential, 115
Counters, 155–156
Cross-sections, 100–101

Degrees of resolution, 97–98
Detours, 16, 213
Diagrams, 3, 28, 67, 81, 83, 105–108, 120, 180
  are not maps, 106–107
  self-referential, 121
Dibutade, 27
Direction of fit, 209
Direction of sight, 103
  oriented, 103
Disorientation, 211–212
Displacement, 32–34, 39
  internal, 24–25

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Distractions, 3, 22, 31, 47, 50, 56, 58, 61, 74, 75, 84–85, 115, 117, 153, 161, 167, 210
  vs. objections, 22
Dot, Green. *See* Green Dot
Dot, Red. *See* Red Dot
Drawings, 26, 36–37, 81
  origins, of 26

Earth's axis, 135
Ecliptic, 134–135
Ecological Psychology, 81
Elgin, Catherine, 168
End-of-Next-Rectilinear-Stretch-Up, 117
Engagement, 211–212
Epiphanies. *See* Representational epiphanies
Epistemic segregation. *See* Segregation, epistemic
Evaluation atoms (formal maps), 51–52
Exclusive Economic Zones, 9
Extractability, 67, 120–121

Flammarion, Camille, 95
Floor, representational, 30–39, 120–121, 147
fMRI, 6, 67, 79, 99
Four Color Theorem, 74–75
Fullness, 177, 179–180

Galilei, Galileo, 32
Generativity, natural, 5, 61–64, 167–168
Geological sequence, 182
GIS software, 79
Globes, 133–144
  faceless, 142
  instructional, 7
  integrated, 135–136, 138, 150
  misoriented, 143–144
  naked, 134–135, 138, 150
  situated, 136–144, 150
  situated, anchoring of, 142–144
  situated, limitations to their use, 140–142
GPS-based navigators, 8, 117
Green Dot, 6, 111 ff., 144, 165, 197

Hippocampus, 214
Holism, 183
Homunculus map, 214
Hutchins, Edwin, 8
Hypergenerativity, 4–5, 63–65, 79, 95–96, 167–168
  specific to maps, 65

Icons, 71
  on maps, 50
Ideographic systems, 174
Illusions, 12
Indexes, 96–98
Indication, 25
  token-reflexive, 152
Inferences, 133–135
  non-situated, 2, 4, 66
  situated, 2, 4, 66–67
  spatial, 79
Informational (non)equivalence, 33, 68
Instruments
  epistemic, 152
  measuring, 152
Intention detection, 29, 96
Intentionality, 96
Islands, 201–208
  advantage, 205
  duality of islands and blocks, 205–206, 208
Itineraries, 53, 174

Jarrett, Keith, 168

Kulvicki, John, 147

Landmarks, 17
Languages, artificial, 183

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Le Carré, John, 193
Light-from-above, bias, 93
Lines on maps, 50
Lists, 24, 52, 173–174
Logic, propositional, 49, 51–52
Lomax, Glen, 161
Looks, 5–6, 89–90
Louis XIV, 130

M1, 2, 12–15, 17–18, 32–35, 37, 39–40, 64, 106, 193, 196
M2, 2, 12–15, 17–19, 32–36, 37–39, 164, 193
M3, 2, 7, 15–18, 39, 164, 196
M4, 2, 7, 15, 17, 24–25, 164
  as proxy mode, 15
Make-believe, 66, 191–192
  navigational, 66
Map layers. *See* Map stages
Map pragmatics, 1
Map stages, 4, 52–54, 70–71, 169–171,
  stages, atomic, 4
Mapmaking, 56, 58, 94
Mappings, 3, 44–46
Maps
  3D, 215
  and territory, 31
  and territory, mutual independence of, 3
  and mappings, 44–46
  and pictures, 83 ff.
  and time, 75–76
  akin to pictures, 4
  a-perspectivality of, 6
  are inference-friendly, 66–67
  are interpreted as factive, 100
  Borgesian (*see* Borgesian maps)
  can lie and distort, 3
  cognition through, 41
  cognitive, 1
  conceptual, 81
  conceptually secondary to GPS, 6
  contrast with language, 54–55
  cortical 81
  create territories, 208–211
  directive, 209
  dynamic, 110, 114–115
  geological, 99
  formal, 4, 5, 49 ff., 55–56, 64–65, 68 ff., 72–73, 75, 79, 86, 169–171
  formal, as the normal form of maps, 55–56
  formal, and mute maps, 56
  formal, expressiveness of, 49, 75
  granularity of, 97–98
  illegitimately taken as pictures, 91
  in the brain, 213–215
  maritime, 31, 76, 118
  mental, 2, 81, 213–214
  mirror-reversed, 72–74
  molecular, 53–54
  navigation friendliness of, 66–67
  not derived historically from pictures, 94–95
  not usually considered as maps, 96 ff.
  of a map, 88
  on the moon (example) 112–114
  ordinary, 68 ff., 79
  ordinary, are abstractions from Borgesian Maps, 128–129
  not ambiguous, 60–63, 177
  not gappy, 60–63, 177
  not glutty, 60–63, 177
  not idle, 60–63, 177
  paper, 211–212
  parts of are semantically active, 56
  picturification of, 5–7, 90, 100, 117
  portability of, 14, 115, 129
  quantized, 97–98, 182
  reshuffled, vii
  retinotopic, 81
  scaled 1:1, 122–123
  self-referential, 121, 124
  situated, 115, 131–132
  specificity of, 9
  superposed, 123, 128

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

syntax, 1
unable to support visual occlusion and perspective, 86
used as models, 133
used for navigation, 18
Marine Protected Areas, 9
Marked arrays, 125
Markings on the ground, 125–127
Mathematical formalism, 35
Memory, 23
and reminiscence, 23
Mental images, 90
as maps, 90
Mereology, 42–43, 53
Mereotopological constraint, 43, 46, 88
Mereotopology, 42–43, 53, 173–174
Metro maps, 53, 81, 189–191
as spatial coordination diagrams, 191
of London, 53
Mind-brain, 11
Mind
"Four-E" account of, 20
augmented, 7
embedded, 23, 114
embodied, 20, 23
enactive, 20
equipped, 7, 19, 23
extended, 2, 7, 19, 23, 114
opportunistic, 19
Misinformation, 159
Misrepresentation, 25, 143–144
Models, 47–48, 133–150
appearance, 134, 144–147, 149–150
architectural, 100
computational, 133–134
concrete, 133–134, 144
mathematical, 133–134
observational, 141
pedagogical, 134
reverse engineering, 134
spatial, 147–149
spatial, different from maps, 149
tailor's, 134
Molecular expressions (music), 175–176, 179
Music formulae, 176–177
Music notation, 1, 48, 167–180
hypergenerative aspects of, 168
naturally generative aspects of, 168
not naturally generative, 167
stems from writing, 167–168
Music pieces, 177
Music scores, 7, 177
formal, 169–171
paced, 177
prescriptive, 168

Nano-cartographies, 98–99
Nano-images, 79, 91
Natural generativity. *See* Generativity, natural
Navigating - vs taking directions, 112
Navigation, 1, 15–18, 39, 109–110, 112, 118, 131
and situation, 110 ff.
as computation, 18
Micronesian, 182–183
possible only if map is situated, 112
spatial, 12
Neo-cortex, 213–214
Non-Overlap, 177, 179–180
Normal forms, 35
North-up, 8, 92–94, 117
Numbered tickets, 197
Numerals, 7, 183–187, 191, 200
sequences of, 1, 182–187

O'Keefe, John, 214
Occlusion shape, 6, 89–90
Occlusion, 85–87, 89–90
Offloading fallacy, 24–25
Offloading, cognitive, 21, 24–25
Ontology, Leibnitian, 72
Ontology, Newtonian, 72
Order, alphabetic, 97

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Organizations, 187–192
fit for mappability, 188
Organizers, 1, 7, 187–192
non-map organizers, 191–192
Orientability, 115
Outline shape, 6, 89–90

Pace, 173
Panoramas, 94, 115–117
Particle colliders, 99
Paths, inferential, 76,78
Perceptual familiarity (perceptual acquaintance), 61–62, 65
Periplii, 53
Perspectival anchors, 118–119
Perspective, 3, 85–86, 117
Photographs, 14, 99, 115–117
aerial, 91–93
situated, 131–132
viewpoint related, 115
Piailug, Mau, 182
Pictograms, 71
Picture Principle (Fodor), 88–89, 179
Picture-based locational algorithms, 99
Pictures, 1, 3, 28, 61–65, 81, 83 ff., 108, 180
are maps of the Albertian Window, 84–85 (*see also* Albertian Window)
are naturally generative, 61–62
are not maps of the scene they depict, 85
are not pictures of the Albertian Window, 85, 89
elicit recognition, 89
hallucinatory, 89
photorealistic, 85
structural accounts of, 87–88
transparency, 88
horse-pictures vs. pictures of a horse, 31
Place cells, 45
Plato, 23
Pliny the Elder, 27
Precomputations, 118
Predication, 69
Principle of Apprehension, 33
Principle of Congruence, 33
Projections, 3, 45–46, 96,
Properties, dissective, 75
Prostheses, 22

Queues, 7, 192–200
as dynamic Borgesian maps, 200
as institutions, 193
as perceptual institutions, 199–200
basic rules of, 193–196, 199
paradox of, 196–199
turn time into space, 194
afford perceptual signals

Radar images, 79
Radars, 6
Random dot stereograms, 89
Rasters, 171–174, 176, 179
as discrete structures, 173
Recordings, 99–101
Recycling, 58–59, 157
neural, 27
Red Dots, 6, 7, 110 ff., 120, 127–128, 142–143, 151–152, 156–158, 160–162, 165–166, 168, 180–182, 187, 190, 192
maintenance, 118
updating, 110–111
Regions, atomic, 97
Representations, 11 ff., 25
Representations
analog, 6, 46–47, 89, 136
hybrid, 67
Representational advantages, 5
Representational epiphanies, 25–28, 39, 61–62, 65, 185
Representational redescription, 27
Representational systems, 29
Rete (astrolabe), 147
Retinal images, 81, 213

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Reynie, Nicholas de la, 130
Rods, measuring, 152–155
Route perspective, 120

Same dimensionality, 41, 47–48, 79, 87, 118, 133, 214–215
Same-dimensionality, principle, 3–4
Satellite images, 92–94
Satellite photography, 105
Satellite views, 5, 66, 83
Schumann Principle, 32–33
Schumann, Robert, 32–33
Sculptures, naked, 134–135
Segments, connected, 172
Segments, initial connected, 172
Segregation, epistemic, 6, 112–115, 131, 142–144, 165
Semantic density, 88
Semantically active zones on maps, 72
Sense data, 5–6, 89–90
Shadows, 7, 27, 91
Shifting, 35
Ship of Theseus puzzle, 28
Shortcuts, 16, 213
Shortening, visual, 101, 103
Shunting, 36–37
Signs, natural, 25–27
Silhouettes, 27
Singular terms, 69–70
  spatially structured, 54–55
Situatedness. *See* Situation
Situation, 6, 109, 115, 127–129, 142–144
  Borgesian maps, 127–129
  as epistemic warrant for map use, 113–114
  different from representation, 109, 115
Sonar images, 79
Sound recordings, 89
Sound reproductions, 79
Southern perspective on mapping, 163
Spelke, Elizabeth S., 186
Staff, musical, 171
Standard Music Notation, 167–180
Star compass, 183, 200
Star names, sequences, 182–183
Star patterns, 7
Stopwatches, 153–156
  different from timers, 153
Stroop effect, spatial, 51
Sundials, 21, 137, 161
  as abstractions from Situated Globes, 140
Symbols, 71
  on maps, 50
  derived, 74–75
  redundant, 74–75
  types of, 68–69
Syncretism, 90
Syntactic sensitivity, 87–88
Systematicity, 41–43, 79, 118, 181

Tabarly, Eric, 58–59
Tables of contents, 96–98
  and indexes, 96–98
Tabula Peutingeriana, 53
Target-up, 117
Team Vestas accident, 8
Territorialization, 2, 6, 208–211
Territories, 3, 6, 7, 9, 31, 47, 59–60, 121, 168–169, 208–211
  independent of maps, 209
Thamus, 23
Thermometers, 81, 192
  do not contain maps, 192
Theuth, 23
Thompson, Nainoa, 182
Ticking mechanisms, 155–156
Time Followers, 152, 155
Time Zones, 159
Timekeeping devices, 151
Timers, 152–156
Timetables, 31–35, 120
Topological entitlement, 86
Topology, 42–43, 53, 86
Tracks, 27

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024

Translating ordinary maps into formal maps, 4, 49 ff., 68 ff., 169–171
  limits to, 50
Trzcinski, Tomasz, 168
Tsivkin, Sanna, 186
Two-system (dual) theories of cognition, 2, 12

Vehicle, representational, 6, 29, 39, 191
Venice, alternate cartography of, 201–208
Verne, Jules, 165
Viewpoint, 117
  and direction of view, 117
  boundedness, 14
  canonical, 6
  located at infinity, 101
Vision, 13–14
Visual angle, 90
Visual routines, 13

Watershed, 30
Wayfinders, deferential, 7
Wayfinding abilities, loss of, 211–212
Web of lines, 153–154
Winds, 79–80
Words, sequences of, 183
Working memory, 24–25
Writing, 1, 3, 7, 27–38, 81, 179
  analogy with music notation, 179
Writing systems, alphabetic, 173–174
Written items are not maps of their content, 174
Written language. *See* Writing
"You are here," 110 ff., 129

Downloaded from http://direct.mit.edu/books/book-pdf/2372891/book_9780262377225.pdf by guest on 18 November 2024